

48048

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Application of : **AFEK et al.**

Serial No.: 09/929,877 : Group Art Unit: 2151

Filed : August 14, 2001 : Examiner: Frantz B. Jean

For : METHODS AND APPARATUS FOR PROTECTING AGAINST
OVERLOAD CONDITIONS ON NODES OF A DISTRIBUTED
NETWORK

Honorable Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

DECLARATION UNDER 37 CFR 1.131

Sir:

We, the undersigned, Yehuda Afek, Anat Bremler-Barr and Dan Touitou, hereby declare as follows:

1) We are the Applicants in the patent application identified above, and are the inventors of the subject matter described and claimed in claims 1-8, 10, 11, 13-16, 20, 33, 35 and 46-69 therein.

2) We conceived our invention prior to September 28, 2000, in Israel, a WTO country. We were then diligent in preparation of a provisional patent application covering the

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

invention during the period between September 28, 2000, and October 17, 2000, when the provisional patent application (US 60/240,899) was filed. The present patent application (US 09/929,877) claims priority from this provisional patent application.

3) As evidence of the conception of the present invention, we attach hereto, as Exhibits A and B, parts of a draft of the present patent application. These documents were prepared September 14, 2000, and September 18, 2000, respectively. (Proof of the dates of these documents, as well as other documents cited herein, is attached hereto as Exhibit G in the form of a directory listing of the archive in which the documents were stored. The relevant files and dates in the archive are noted below.)

4) The following tables show the correspondence between the independent claims now pending in this application and Exhibits A and B. In view of this correspondence, it is clear that we conceived the claimed invention prior to September 28, 2000.

Claim 1	Exhibits
A method of responding to an overload condition at a network element ("victim") in a set of one or more potential victims on a network	Exhibit A, page 1, paragraph 1: "NetGuard system is activated upon receiving alerts of an attack. The system then focused on defending only the victim(s) of the attack."

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

A. responsively to an indication of an anomalous traffic condition, initiating diversion of traffic destined for the victim by a first set of one or more network elements external to the set of one or more potential victims to a second set of one or more network elements external to the set of one or more potential victims	Exhibit A, page 1, paragraph 4: "At the time of the attack all traffic to the server, which is the victim of the attack, is navigated to the NetGuard. This is done by routing any traffic using the <i>victim public address</i> to NetGuards. Hence achieving our first goal, that traffic to the victim, from outside the network, and inside the network, is redirected to NetGuards."
B. the element(s) of the second set filtering traffic diverted in step A ("diverted traffic") and selectively passing a portion thereof to the victim.	Exhibit A, page 1, last paragraph: "The NetGuards machine, discriminates between traffic to the victim that is part of the attack, and genuine traffic. The traffic of the attack would be blocked at NetGuards. Genuine traffic would be routed from the NetGuards to the victim, using the <i>victim private address</i> ."

Claim 46	Exhibits
A network element for use in protecting against an overload condition on a network	Exhibit A, page 1, paragraph 1: "NetGuard system is activated upon receiving alerts of an attack. The system than focused on defending only the victim(s) of the attack."

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

<p>an input for receiving traffic diverted from the network, the traffic comprising flows of data packets having respective source addresses</p>	<p>Exhibit A, page 1, paragraph 4: "At the time of the attack all traffic to the server, which is the victim of the attack, is navigated to the NetGuard. This is done by routing any traffic using the <i>victim public address</i> to NetGuards."</p> <p>Exhibit B, section 1.1: "It is common (e.g., in the Cisco convention) to define a network flow by the following parameters: i. Source IP address..."</p>
<p>a statistics module that is arranged to perform a statistical analysis of the diverted traffic so as to detect an anomalous pattern of a flow associated with at least one of the source addresses</p>	<p>Exhibit B, section 1.3.2: "Attack Analysis: Will be conducted during attack time and will be responsible to compare the historically collected statistical data with the current traffic volume and generate rules for traffic blockage. The output of this unit, in general, will consist of a list of items for each of which three parameters will be provided: a. Network flow, identified by a combination of source IP address (can be prefixed), destination IP address, destination port number, protocol type..."</p>

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

a filter, which is operative, responsively to detection of the anomalous pattern, to block at least a portion of the data packets having the at least one of the source addresses	Exhibit B, section 1.3, last paragraph: "The analysis will be based on the statistical parameters of the data and will aim at keeping the attacked destination at normal loads by blocking the most 'suspected' traffic streams."
an output coupled to the input for selectively passing on to further elements in the network traffic not blocked by the filter	Exhibit A, page 1, last paragraph: "The NetGuards machine, discriminates between traffic to the victim that is part of the attack, and genuine traffic. The traffic of the attack would be blocked at NetGuards. Genuine traffic would be routed from the NetGuards to the victim, using the <i>victim private address</i> ."

Claim 46	Exhibits
A system for use in protecting against an overload condition on a network	Exhibit A, page 1, paragraph 1: "NetGuard system is activated upon receiving alerts of an attack. The system then focused on defending only the victim(s) of the attack."
one or more network elements ("guards") disposed on the network	Exhibit A, page 1, paragraph 4: "At the time of the attack all traffic to the server, which is the victim of the attack, is navigated to the NetGuard."

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

an input for receiving traffic from the network	Exhibit A, page 1, paragraph 4: "This is done by routing any traffic using the <i>victim public address</i> to NetGuards."
a filter coupled to the input, the filter selectively blocking traffic originating from a source suspected as potentially causing the overload condition	Exhibit B, section 1.3, last paragraph: "The analysis ... will aim at keeping the attacked destination at normal loads by blocking the most 'suspected' traffic streams."
a statistics module that is coupled to the filter and that identifies the traffic statistically indicative of having originated from the source suspected as potentially causing the overload condition	Exhibit B, section 1.3.2: "Attack Analysis: Will be conducted during attack time and will be responsible to compare the historically collected statistical data with the current traffic volume and generate rules for traffic blockage. The output of this unit, in general, will consist of a list of items for each of which three parameters will be provided: a. Network flow, identified by a combination of source IP address (can be prefixed), destination IP address, destination port number, protocol type..."

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

an output coupled to the input for selectively passing on to further elements in the network traffic not blocked by the filter	Exhibit A, page 1, last paragraph: "The NetGuards machine, discriminates between traffic to the victim that is part of the attack, and genuine traffic. The traffic of the attack would be blocked at NetGuards. Genuine traffic would be routed from the NetGuards to the victim, using the <i>victim private address</i> ."
one or more further network elements ("diverters") disposed on the network and in communication with the guards, the further network elements selectively initiating, responsively to detection of an anomalous traffic condition, diversion to at least one of the guards traffic otherwise destined for a still further network element ("victim") in a set of one or more potential victims on the network	Exhibit A, page 1, "routers" shown in the figure diverting traffic to "NetGuards," as stated in paragraph 4 on page 1: "At the time of the attack all traffic to the server, which is the victim of the attack, is navigated to the NetGuard. This is done by routing any traffic using the <i>victim public address</i> to NetGuards. Hence achieving our first goal, that traffic to the victim, from outside the network, and inside the network, is redirected to NetGuards."

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

Claim 56	Exhibits
A method of responding to an overload condition at a network element ("victim") in a set of one or more potential victims on a network	Exhibit A, page 1, paragraph 1: "NetGuard system is activated upon receiving alerts of an attack. The system then focused on defending only the victim(s) of the attack."
diverting to a guard machine traffic destined for the victim, the traffic comprising flows of data packets having respective source addresses	Exhibit A, page 1, paragraph 4: "At the time of the attack all traffic to the server, which is the victim of the attack, is navigated to the NetGuard. This is done by routing any traffic using the <i>victim public address</i> to NetGuards." Exhibit B, section 1.1: "It is common (e.g., in the Cisco convention) to define a network flow by the following parameters: ii. Source IP address..."

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

performing a statistical analysis of the diverted traffic at the guard machine so as to detect an anomalous pattern of a flow associated with at least one of the source addresses	Exhibit B, section 1.3.2: "Attack Analysis: Will be conducted during attack time and will be responsible to compare the historically collected statistical data with the current traffic volume and generate rules for traffic blockage. The output of this unit, in general, will consist of a list of items for each of which three parameters will be provided: a. Network flow , identified by a combination of source IP address (can be prefixed), destination IP address, destination port number, protocol type..."
a filter, which is operative, responsively to detection of the anomalous pattern, to block at least a portion of the data packets having the at least one of the source addresses	Exhibit B, section 1.3, last paragraph: "The analysis will be based on the statistical parameters of the data and will aim at keeping the attacked destination at normal loads by blocking the most 'suspected' traffic streams."

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

responsively to detecting the anomalous pattern, preventing at least a portion of the data packets having the at least one of the source addresses from reaching the victim while passing to the victim at least some of the data packets from other source addresses	Exhibit A, page 1, last paragraph: "The NetGuards machine, discriminates between traffic to the victim that is part of the attack, and genuine traffic. The traffic of the attack would be blocked at NetGuards. Genuine traffic would be routed from the NetGuards to the victim, using the <i>victim private address</i> ."
---	---

Claim 66	Exhibits
A method of responding to an overload condition at a network element ("victim") in a set of one or more potential victims on a network	Exhibit A, page 1, paragraph 1: "NetGuard system is activated upon receiving alerts of an attack. The system than focused on defending only the victim(s) of the attack."
coupling the victim to receive traffic from the network via a first port of a network switch	Exhibit A, page 1: In the figure, the victim is coupled to receive traffic via one output of a router.

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

actuating the network switch to divert the traffic destined for the victim to a second port to which a guard machine is coupled	Exhibit A, page 1, paragraph 4: "At the time of the attack all traffic to the server, which is the victim of the attack, is navigated to the NetGuard. This is done by routing any traffic using the <i>victim public address</i> to NetGuards." The figure shows that the NetGuard is coupled to a different port of the router from the victim.
filtering the diverted traffic using the guard machine	Exhibit A, page 1, last paragraph: "The NetGuards machine, discriminates between traffic to the victim that is part of the attack, and genuine traffic. The traffic of the attack would be blocked at NetGuards."
selectively passing at least a portion of the filtered traffic from the guard machine to the victim	Exhibit A, page 1, last paragraph: "The traffic of the attack would be blocked at NetGuards. Genuine traffic would be routed from the NetGuards to the victim, using the <i>victim private address</i> ."

5) During the period between September 28 and October 17, we worked continuously and diligently to revise and supplement the material in the original drafts in order to complete the provisional patent application that was subsequently filed. Some of the draft documents that we prepared during this period are attached hereto as Exhibits C, D, E and F. These documents were completed, respectively, on September 29, October 2, October 9, and October 13, 2000. We then

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

completed and filed our provisional patent application on October 17, 2000.

6) Exhibit G is a directory listing of the archive from which Exhibits A-F were taken. The table below lists the file names and dates as they appear in Exhibit G:

Exhibit	File Name	Date
A	Netxxn.doc	September 14, 2000
B	Statistical-patent4.doc	September 18, 2000
C	Copy of netxx.doc	September 29, 2000
D	Attack Identification.doc	October 2, 2000
E	Statistical-patent-hanoch5	October 9, 2000
F	Mordi.ppt	October 13, 2000

US 09/929,877

Declaration under 37 C.F.R 1.131 by Afek et al.

We hereby declare that all statements made herein of our own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application of any patent issued thereon.

Yehuda Afek
Citizen of Israel
26 Hacarmel Street
Hod Hasharon
Israel

Date:

9/10/08

Dan Touitou
Citizen of Israel
21 Golani Street
Ramat Gan 52224

Israel:

Date: 9/10/08

Anat Bremler-Barr
Citizen of Israel
17 Hashomron Street
Ramat Hasharon
Israel

Date:

1. Activation netGuards system

As described above, NetGuard system is activated upon receiving alerts of an attack. The system then focused on defending only the victim(s) of the attack. In the first version of the NetGuard system we assume that the information about the existing of the attack, and the information who is the victim is injected to the system from outside.

Activation the NetGuard system enforces two important change in the flow of traffic to the victim:

1. Traffic to the victim, from outside the network, and inside the network, is redirected to NetGuards.
 2. Only flow that passes throw NetGuard can reach the actual victim.
- In this section we describe in details one of the possible architecture and mechanism that achieve the above goals.

We give to each server IP two IP addresses. One is the server *public address* and the other is the *server private address*. The *server public address* is the address of the server that is spread in the world, throw the DNS mechanism. The *server private address* is the address that known only to trustable parts of the networks, i.e., the NetGuards and to the interfaces of routers that connected to routers or netGuards (See figure 1). In other words, the *server private address* is not known to router interfaces that are connected to hosts. This give us the ability, to discard packets originated from hosts, that uses the *server private address*.

At the time of the attack all traffic to the server, which is the victim of the attack, is navigated to the NetGuard. This is done by routing any traffic using the *victim public address* to NetGuards. Hence achieving our first goal, that traffic to the victim, from outside the network, and inside the network, is redirected to NetGuards. We give below details of one of the possible ways to i redirect the traffic (see subsection 1.1)




Figure 1: The *victim private address* is known only to trustable parts in the networks, i.e., the interfaces of routers that connected to another router or to netGuards. The routes in the network where the *victim private address* is known is marked by dashed red lines.

The NetGuards machine, discriminates between traffic to the victim that is part of the attack, and genuine traffic. The traffic of the attack would be blocked at NetGuards. Genuine traffic would be routed from the NetGuards to the victim, using the *victim*

private address. This would be done by a simple manipulation on the destination field in the packets of the genuine traffic. Hence we achieve our second goal, and only flow that passes through NetGuard can reach the actual victim. This is due, to the fact that in our architecture only traffic originated from NetGuard can use the *victim private address*, and hence reach the victim.

When the attack is ended, we just cancel the redirection to NetGuards, of all the traffic to the victim. Hence, again any traffic addressed *victim public address* can be routed directly to the server, which was the victim of the attack. In the following subsection we give the details of the redirection mechanism.

1.1 Redirecting traffic to NetGuards

Recall, that in case of attack we want to route any traffic addressed to the *victim public address* to NetGuards. This can be done by two different methods, one more suitable for the traffic that originated inside the network and the other more suitable to the traffic that originated outside the network.

Traffic that originated outside the network, would surely need to pass through one of the border routers, i.e., a router that surrounded the network. In our architecture a NetGuard machine is attached to each such border router. In time of attack, the control of the NetGuards system injected an update to the border router, by updating the policy routing mechanism in the router. This update, would notify the border router that any packets that is addressed to the *victim public address* would be forwarded to the NetGuards. Updating the policy routing mechanism, gives us the ability to change the routing behavior without degrading the border routing performance¹.

In case the traffic, is originated from inside the network, one could use the same mechanism, in order to redirect the traffic addressed to the victim to the NetGuard. However, this required updating the policy routing of all the routers in the network. Hence in many cases, it is more beneficial to use a different method, that is based on a simple routing manipulation. The designated NetGuard² that handles the inside traffic to the victim would announce its IP address as the *victim public address*, while the

¹ Page: 2 Unlike access list, that required filtering every packet, and hence degrading the router performance, Policy routing does not harm the routing performance.

To understand this, we briefly explain the look up process in today routers. Most of the routers use Cisco Express Forwarding, or some equivalence mechanism. Using FEC, every interface has a cache where it stores the information about the next hop for the last packets that arrive through this interface. When packets arrive to the interface card and the destination is not stored in the cache, a new forwarding process is done. This process is done in the central unit that does the lookup process for all the interfaces. This process takes into account the routing policy that can be defined per interface. This operation is done rarely and in almost all of cases, the lookup operation uses the cache information. Hence the impact of the degrading in the forwarding time is minor.

² In some case in order to handle the volume of the intra network traffic, it may be beneficial to use not one NetGuard, but a farm of NetGuards. However, one should notice that the problem of attacks, and special spoofing attack, in most of the cases is harder when the attack is originated from outside the network. When the attack is originated from inside network, there is full information and management of the network. Hence ingress and egress filtering can be used, for dealing with spoofing attack. In cases when the origins of the attack are known, one can more easily stop the attack, by disabling the origins of the attack.

EXHIBIT A

victim server, would redraw from this address. This routing updating information, would spread quickly in the networks, using the standard routing protocol, e.g., OSPF, EIGRP or RIP.

TCP Anti-spoofing

We describe here a new anti-spoofing techniques which is TCP oriented. This anti-spoofing mechanism authenticated the genuine of the source address of the flow, based on the SYN mechanism of TCP. When a host wants to open a TCP connection with the server, the hosts sends a SYN request, notified about its wish for a new connection. The server authenticated its source address by sending him back a random number. Then, the server wait to received this number back the source. Naturally a spoofed source cannot repeat the number, and hence any connection between the server and a spoof source, is dismissed. Hence the SYN phase, which is the connection establishment phase in TCP (also called the three-way handshake), is a naturally anti-spoofing method (see figure 2).

However, since this mechanism is done by the server, the SYN mechanism has become one of the efficient way to do denial of service attack. SYN-attack, is based on the fact that the server get high volume of SYN request. This lead to the fact that the buffer, of SYN requests is filled, dismissing any new SYN requests, which can be a SYN request of a genuine host. This kind of attack also make a huge burden on the CPU server.

In our architecture we built a purposed computer, the NetGuard computers, that take the role of the server and do the SYN process. The NetGuards, can deal with a high volume of traffic, which in many cases equal to the bandwidth of the links. The architecture of the NetGuards system also distributed the load of the attack in the SYN-attack, on the number entrance points to the network.

Using a special computer for the SYN process, is very naturally solution. Also in day to day life, the job of guarding and gatekeeping is separated from the actual activity that is guarded due to performance issue.

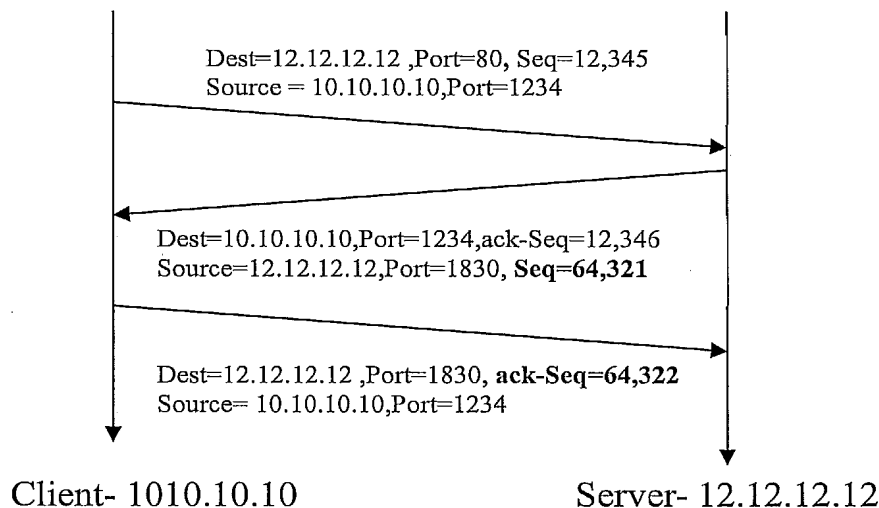
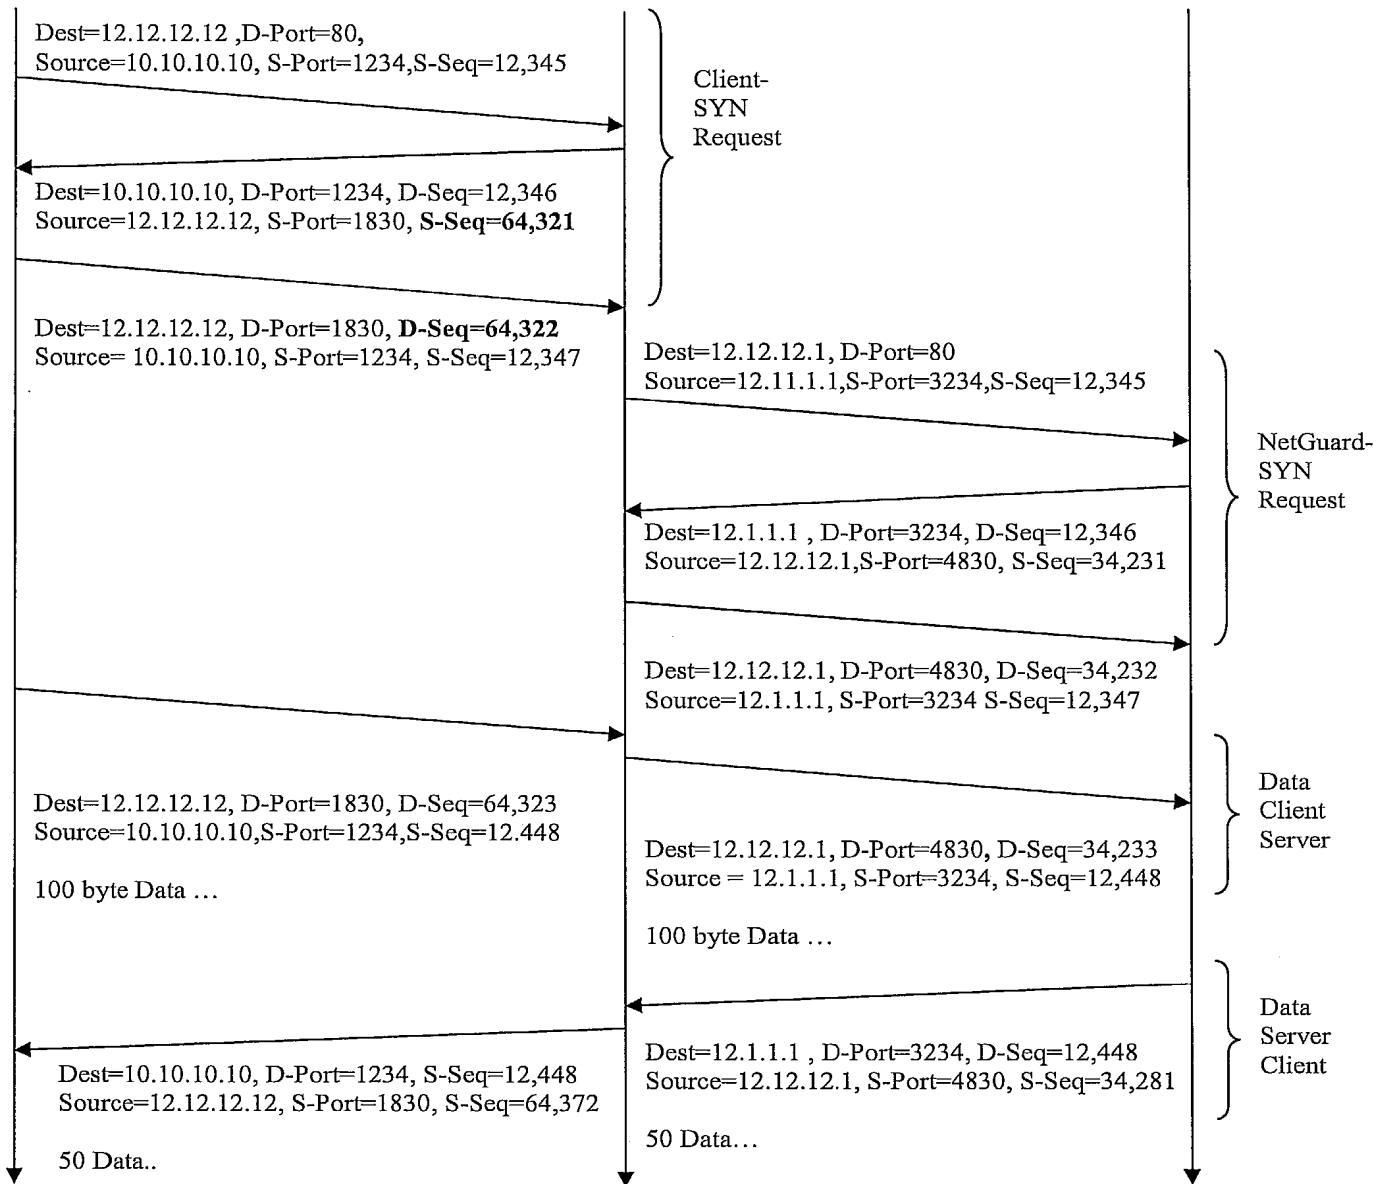


Fig 2: The SYN request.

EXHIBIT A



Client- 10.10.10.10

NetGuard - 12.1.1.1
Play the rule of 12.12.12.12
12.12.12.12 is the server
public address

Server- 12.12.12.1
this address is the
private address of the
victim

EXHIBIT A

, filling the buffer of the server, with spoofed

many spoofed source start the operation of the SYN-attack

If the source addressed is gennot spoof, than The basic idea is to send back to the source a random number.

Authenticated the source of the flow. Thus distinguish between spoof source address to real source address. The Authentication mechanism is a new anti-spoofing techniques TCP oriented

Ability to throw up -

To divide the work of the routers...

1. Attack Identification, Recognition and Isolation via the Statistical Recognition Unit.

The statistical recognition unit is responsible for analyzing the attack, identifying its origin(s) and providing operational rules for blocking the attack without disturbing innocent genuine traffic. The basic principle behind the unit's operation is that the pattern of traffic originated at the attack sources during attack time drastically differs from that pattern during normal operation. In contrast, traffic patterns of "innocent" sources during an attack resemble those at normal times. This principle is used to identify the attack sources and provide guidelines for their blockage.

The statistical unit has three major components: a) Classification of the victim traffic into flows b) Learning the traffic patterns of the various victim flows under normal operation conditions, and c) Monitoring the flows traffic pattern at attack time and detecting the attack sources. Below we describe these in detail.

1.1 Network flows and traffic classification

The statistical unit operation is based on classifying the traffic into network flows. A network flow can be viewed as a stream of packets that share the same properties. It is common (e.g., in the Cisco convention) to define a network flow by the following parameters:

- i. Source IP address.
- ii. Source port.
- iii. Destination IP address.
- iv. Destination port.
- v. Traffic type (TCP / UDP / SYN).

ZZZ will use either this fine classification or a more coarse classification, guided by the following considerations:

- i. **Disregarding source port:** Will be done in the event that source port does not serve as a good separator between malicious and innocent traffic.
- ii. **Grouping (aggregating) a set of individual source addresses into one set (e.g., by considering the IP address prefix):** Aggregation, if used, will serve to reduce the number of statistics measured and computed at attack time, thus reducing the processing complexity. Aggregation can be done in a hierarchical manner.
- iii. **Disregarding destination address:** At most cases the unit operates to block attacks oriented on a single target (e.g. www.xyz.com). In these cases, the statistical unit will receive as input *only* flows destined to that destination (www.xyz.com). In these cases classifying by destination address is irrelevant.

1.2 Traffic Studying During Peace Time

During “peace time” the unit will actively measure and study the traffic volumes of the various flows. This is done in two major modules:

1. **Traffic volume statistical data collection and classification:** This module operates at “peace” times and is destined to learn the normal traffic volume patterns. This traffic learning is done in either (or both) of the following approaches:
 - i. Sampling a **fraction** α of the **packets** ($0 < \alpha \leq 1$) traversing the lines on route to the target and then classifying the packets based on the following parameters:
 - x. **Network flow** – classification according to the classification described in the previous sub-section.
 - y. **Time of the day and day of the week.**

Note that setting $\alpha=1$ requires the unit to process every packet and thus imposes high load on it while providing the best statistical measure. Lower values of α reduce the load posed on the unit while potentially somewhat degrading the statistical measure. The fraction α , therefore, will be a parameter that will be set so that enough statistical knowledge can be gained without over-loading the system.

This method can be used to classify all traffic type direct to/from the defended targets, and requires sensing (“sniffing”) the lines on route to the destination. The sniffing devices must be placed as to measure all traffic, that is, at the network boundary, or at the defendant target proximity.

- ii. Utilizing **server logs** collected by the defended target. These typically contain information about the activity being performed on the target. For example, WEB sites, which are likely to form the main body of potential targets, keep logs that record all the document requests sent to the site (including their source address, time of the day and other parameters). Processing of these logs by ZZZ will yield a very accurate measure of the statistics of network flow volumes (measured in packets per second, as in a) above).

The traffic volume data collected will be summarized and will be stored in a database that can be accessed via the various parameters of the flows.

2. **Traffic analysis:** Will be conducted at “peace time” and used to generate statistical summaries of the data collected. In particular the processing will be used to compute **mean** and **variance** of volumes of

EXHIBIT B

each of the flows, or aggregates of flows. The analysis may also dynamically change aggregates of flows in order to improve the statistical identification of traffic. The results of this analysis will be stored in a database to be used at attack time.

1.3 Traffic Measurement and Analysis at Attack Time

1. Online traffic volume collection at attack time: This module operates during attack times and is responsible to collect the statistics of the traffic at that period. The module receives as input only traffic that is destined to the attacked target(s) and measures its packet rates. Note, that in this sense, its measures are similar to the “peace time” measures collected in approach 1a above. The classification of the traffic, in general, is similar to that conducted in “peace time” but may be controlled/guided by external intervention. Such intervention will be enacted if some additional knowledge on the attack type is gained from other sources (e.g., human-aided identification) and can be utilized by the unit.

2. Attack Analysis: Will be conducted during attack time and will be responsible to compare the historically collected statistical data with the current traffic volume and generate rules for traffic blockage. The output of this unit, in general, will consist of a list of items for each of which three parameters will be provided:

- Network flow**, identified by a combination of source IP address (can be prefixed), destination IP address, destination port number, protocol type.
- Duration**, identifying the duration for which that class will be blocked.

The analysis will be based on the statistical parameters of the data and will aim at keeping the attacked destination at normal loads by blocking the most “suspected” traffic streams. Blocking rules will be based on maximizing the likelihood of blocking malicious data while minimizing the likelihood of blocking innocent data.

1.4 Statistical Recognition of Data “Innocence”

ZZZ will use two major properties of network flows to identify whether they are malicious or innocent. These are: a) Traffic pattern, and b) Traffic volume. Below we describe the recognition approaches based on these factors.

1.4.1 Recognition of Traffic Pattern

Several aspects of traffic pattern will be examined:

- 1) Source “IP geography” proximity:** Sources will be classified into classes that resemble the “IP geography”, that is IP addresses that reside on similar networks (similar IP address prefix) will be classified in the same class. A class that will generate a relatively-large volume of requests will be suspected as being malicious. Note that such “malicious classes” are likely to form if the attacker planted a collection of daemons in the same network.

- 2) **Periodicity:** Sources will be examined for the periodicity of their requests. It is likely that malicious sources (unless being very sophisticated) will act in a relatively periodic manner, while innocent sources act in more random fashion.
- 3) **Packet Properties:** Sources will be examined for repetitive properties of their packets. For example the distribution of packet size. It is likely that malicious sources (unless very sophisticated) will generate packets of identical properties (e.g. – all packets of same size) while innocent sources will generate packets of more random nature.

1.4.2 Recognition of Traffic Volume

Traffic volume recognition will be used to identify malicious sources that transmit **large volumes** of data which **significantly differ** from their normal volume. Specifically, we classify Internet data sources to *small sources* and *large sources*. The former relates to individual IP addresses whose traffic volume is normally tiny. The latter relates to Proxy traffic or Spider traffic¹ whose volume is drastically higher.

ZZZ will keep individual volume history for each of the large sources. Individual history will not be kept for the small sources; rather a single fixed small number (related to their mean volume averaged over all these sources) will be recorded. At attack time the traffic volumes of individual flows will be measured and compared to their recorded volume. Flows whose volume will drastically differ (upwards) from their recorded measure will be marked as being malicious.

The mathematical formulation of this procedure is as follows: Given are K classes of flow of traffic, indexed $1, 2, \dots, K$, and characterized by the mean (μ_i) and the variance (σ_i) of their historical volume, and by their current volume (X_i). We would like to identify the classes which mostly deviate from their expected volume. Let $Y_i = (X_i - \mu_i) / \sigma_i$. We will sort the classes by the value of Y_i and will recommend to block (eliminate) the classes with the largest values of Y_i .

1.4.2.1 Time accumulating traffic volume recognition and “controlled” denial of service

It is important to recognize that the effectiveness of volume recognition increases with the time duration along which it is implemented. This is correct since the variance of total data volume generated by a source during a period of duration T decreases in T . For example, it is expected that the average amount of traffic generated by a small source during a period of 1 hour will be *very small*. However, at certain epochs, it is expected that the average amount of traffic generated by the same source during a period of 1 minute can be rather large. (up to 60 times larger than that of the 1 hour average).

¹ The traffic volume resulting from a Spider access is normally higher than that of a single human user, especially on large WEB sites. The reason is that a spider scans the whole site, leading to hundred or thousands of requests while a human client requests tens of pages or less, on average.

EXHIBIT B

For this reason ZZZ will implement the following unique recognition and traffic screening mechanism. For source i , let $S_i(t)$ denote the amount traffic generated by the source during the interval $(0, t)$ (where we assume that the attack starts at time 0). We then set at time t : $X_i(t) = S_i(t)/t$ and apply the above screening mechanism.

This mechanism has the following properties:

1. For a small value of t (that is, at the attack beginning moments) a sophisticated attacker might cause significant number of innocent users denial of service. This is due to the fact that the attacker may inflict a load that resembles that of an innocent client, and thus the attacker is not distinguishable from the innocent client. At this stage, ZZZ will may block some innocent clients and some attackers. Using this action, for a short period some innocent clients may be denied of service but ZZZ protects the site from going down!
2. As t increases the denial of service conducted by ZZZ will be acted more and more on the malicious sources and less and less on the innocent sources. This is since the malicious sources have posed large amount of accumulated load. Thus as time progresses less and less innocent clients are denied of service. In fact, after relatively short period all malicious sources will be denied of service while the innocent sources will receive full regular service.

Example: Consider the traffic volume generated on the web site of the Nagano server (Feb 98). It had 11,665,713 requests made over a period of 24 hours by 59,582 clients. Assuming uniform distribution of clients over the day, this implies about 2500 clients per hour and 500 clients per 12-minute interval. An attacker who uses 500 sophisticated daemons (which imitate a normal client) will look innocent at the first 12 minutes interval. At this period ZZZ will block 50% of the innocent clients and 50% of the daemons. However, after 24 minutes the daemons will generate significantly more traffic than an innocent client and thus almost all of the traffic blocked will be that of daemons.

1. Abstract

	DISTRIBUTED NETWORK DEFENSE SYSTEM	1
1	ABSTRACT	1
2	INTRODUCTION	1
3	ACTIVATING THE GUARDS SYSTEM	3
	3.1 Redirecting traffic to NetGuards	3
4	TCP ANTI-SPOOFING	6
5	INGRESS FILTERING	9
6	ATTACK IDENTIFICATION, RECOGNITION AND ISOLATION VIA THE STATISTICAL RECOGNITION UNIT.	9
	5.1 Network flows and traffic classification	9
	5.2 Learning Traffic Characteristics	10
	5.3 Traffic Monitoring and Analysis at Attack Time	11
	5.3.1 Statistical Recognition of Data "Innocence"	12
	5.3.1.1 Recognition of Traffic Pattern	12
	5.3.1.2 Recognition of Traffic Volume	12
	5.3.2 accumulating traffic volume recognition and "controlled" denial of service	Time 13

2. Introduction

Despite numerous distributed denial of service attacks that took place last year (with a surge of attacks on YAHOO, CNN, and many other major sites), there is still no known online solution that directly protects during an attack. Here we present a distributed defense system that does this, enabling the continuous operation of an attacked site while the attack is going on.

Our method of protecting victim elements in the network during an attack and enabling them continuous operation despite the attack follows these major steps:

a. **Detection:** We assume that an attacked victim has an automatic mechanism that detects and alerts when an attack begins. Such a mechanism is provided by different routers, firewall equipment and operating systems.

b. **Alert:** Upon suspecting that an attack has began an alert network activates a network of guard machines that are located at strategic points around an area of the network in which the victim resides. For example, the guards could be located around the autonomous system that hosts the victim.

c. **Divert:** In addition the alert network invokes a rerouting mechanism that ensures first that all the traffic destined to the victim is diverted to the guards and second that no packet (message) reaches the victim unless it passed through a guard.

d. **Sieve:** When an alert arrives at a guard and the victim traffic is being received, the guard sieves the traffic to sort out the “bad” packets and pass on to the victim only the “good” traffic. The architecture and operation of the special purpose guard machine is the second part of this patent and has the following four major components (see Figure 1):

1. **Anti-spoofing:** An anti-spoofing module that authenticates and verifies for each flow (<source-address, source-port, destination-port> triplet) that a real process at the host with that source-address and behind that port-number has initiated this flow.

2. **Statistics:** A module that detects and singles out flows (Source IP addresses or subnetworks) with outstanding behavior. The identity of these flows is passed to a filter.

3. **Filter:** A module that blocks any packet originating from an IP address or subnetwork that was identified in the previous step as a source of malicious traffic.

4. **Ingress filtering:** The guard machines interact with its neighboring routers to enable an effective usage of the ingress filtering feature. The routers do not always know which flows they may block because of route asymmetry. The guards analysis of the traffic both at normal operation and during an attack would to pin point which IP addresses and addresses blocks may be purged.

5. **Termination detection:** All the guards participate in a fourth module which is a distributed algorithm they run to cooperatively decide when an attack has stopped and the victim may return to peaceful operation mode. This last transition has to hand over the good connections from the guards to the victim.

Packets flow through a guard machine by first passing through the third component, then through the first and finally through the statistical module.

EXHIBIT C

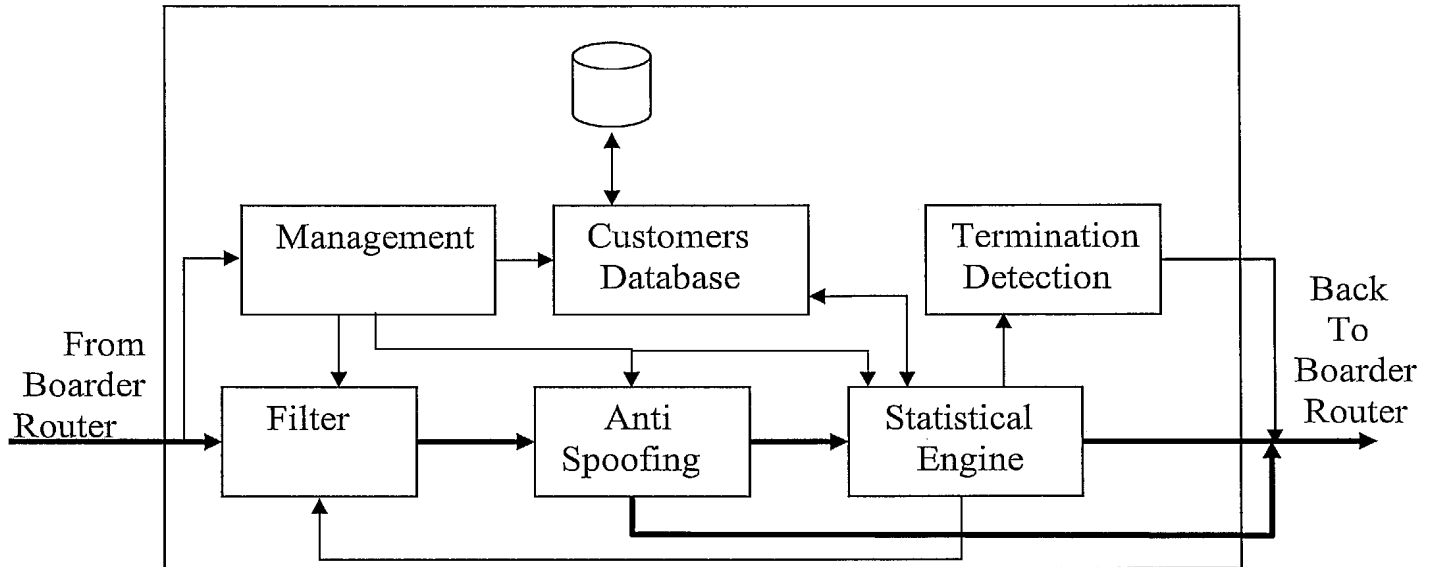


Figure 1: High level block diagram of a guard machine. Thick lines represent the flows of victim traffic through the machine while the thinner lines represent control and other information paths. Customers database maintains statistical information about each potential victim traffic patterns (each customer is a potential victim that the guard would protect).

In addition we describe how to deal with the open connections while transiting from normal operation to protect mode, and how to deal with the open connections while transiting from protect mode back to normal mode (after an attack has been stopped).

The diversion method that passes all the victim traffic, during an attack, through the guard machines is described in the next section. In Section 4 the antispoofing module is described and the

3 Activating the Guards system

3.1 Redirecting traffic to NetGuards

Upon suspecting that an attack is being mounted on it, a victim alerts the guards through a communication channel supplied by the backbone provider. For example by sending authenticated messages to the NOC (Network Operations Center) from which the message is relayed to the guards (SNMP and other network management means may be used instead). The alert message contains the identity of the victim machines (which includes their IP addresses). At this point the victim enters the "protected" mode.

In this "protected" mode all the traffic flows to the victim have to be diverted such that:

1. All the traffic whose destination is the victim, from either outside the autonomous system (AS) that hosts the victim or from inside the AS, is redirected to the guards.

2. All the traffic that reaches the victim must have passed through one of the guards, i.e., no traffic can reach the victim without going through the guards.

In this section we describe one possible architecture and mechanism to achieve the above goals.

Two IP addresses are associated with each victim destination machine (server) the *server (victim) public address* and the *server (victim) private address*. The *server public address* is the IP address of the server which is published all over the Internet through the DNS system. The *server private address* is an IP address used solely to transfer packets between the guards and the victim, while the guards protect the victim. Therefore, the *server private address* is recognized only by either router interfaces that are connected to internal links of the hosting AS backbone, or to the guards, or to the victim interface. All other interfaces, such as, connections to links that come from outside the AS, or links that are connected to other customers of the AS (stub networks), discard any packet whose destination is the *server private IP address* (easily and efficiently achieved by using the CEF mechanism of the routers). (See figure 1). This ensures that no hacked daemon can generate traffic to the *victim private IP address*.

To achieve the above setting all the interfaces that are connected to external links, i.e., links that connect to either other networks (AS's) or to external hosts and customer networks, are permanently programmed to discard traffic destined to the *server private IP address*. In normal operation when no attack is being mounted, the victim declares itself to be at distance zero from both the *server private IP address* and the *server public IP address*. This causes the routing protocol to set entries in the forwarding tables in all the AS routers, to forward messages destined to either address to the victim (which is now not a victim) machines.

To divert *public IP address* victim traffic that arrives from outside the hosting AS during an attack we notice that all such traffic must pass through one of the border routers, i.e., a peering or NAP, BGP routers. A guard machine is placed in each entry next to the boarder routers at this point. Upon receiving the alert of a possible attack on a victim all these boarder routers are set to forward all the traffic arriving from out of the network (AS) and whose destination address is the *victim public IP address*, to the guard machine which is placed next to them. This is easily achieved by injecting an update to the boarder routers, updating their policy routing mechanism. Updating the policy routing mechanism, gives us the ability to change the routing behavior without degrading the border routing performance¹. In effect the guarding machines

¹ Unlike access list, that required filtering every packet, and hence degrading the router performance, Policy routing doesnot harm the routing performance.

To understand this, we briefly explain the look up process in today routers. Most of the routers use Cisco Express Forwarding, or some equivalence mechanism. Using FEC, every interface has a cache where it store the information about the next hop for the last packets that arrive throw this interface. When packets arrive to the interface card and the destination is not store in the cache, a new forwarding process is done. This process is done in the central unit that does the lookup process for all the interfaces. This process take into account the routing policy that can be defined per interface. This operation is done rarely and in almost all of cases, the lookup operation use the cache information. Hence the impact of the degrading in the forwarding time is minor.

become a TCP proxy for the victim. All the traffic returning from the victim to trusted clients is passed through the corresponding guarding machine.

To divert victim *public IP address* traffic that originates inside the hosting AS to the internal or boarder guards, one could use a similar mechanism. That is, to inject when the alert is received, the desired routing information into all the routers. However, this requires updating the policy routing of all the routers in the AS. In many cases (large networks), it is more beneficial to use a different method, based on a simple routing manipulation. In this method, when the victim suspects that an attack is being applied, it declares itself to be at a large distance from the *server (victim) public IP address*, while the guards would start to declare that they are at distance zero (or close to zero) from the *server (victim) public IP address*.² This routing updates quickly spread in the AS network, using the standard routing protocol (usually a link state type of protocol), e.g., OSPF, EIGRP or RIP³. Thus within seconds from these declarations all the victim traffic is automatically diverted to the guards.

When the guards decide (see how below) that the attack has terminated they send an appropriate message to the victim machine. At the same time they reverse the above settings, that is, they stop declaring that their distance from the *server (victim) public IP address* is zero, while the victim starts declaring again that it is at distance zero from its public IP address.

Notice that in the “protect” mode several guards may all claim to be at distance zero from the *victim public IP address*. This divides the AS into clusters, such that packets with this destination address in each cluster are routed to the guard residing within that cluster. However, there might be routers on the boarder between two or more such clusters with equal distance to two or more guards. This may introduce routing instability, where some packets of a flow go to one guard and some packets go to the other guard. First notice that this effects only victim traffic that originates inside the hosting AS. The victim traffic that arrives from outside the hosting AS is treated by the guard at the entry point which acts as a proxy for that traffic. Thus outside traffic would suffer from route flapping only if these flapping are introduced by BGP, which is very rare. To avoid the flapping of victim traffic originating inside the AS, we set each guard to declare that it is at a very small but different distance from the victim *public IP address*. This small perturbations ensure that no router would be at equal distance from two guards. The exact calculation of this perturbation is automatically calculated given the ISP map of its backbone.

² In some case in order to handle the volume of the intra network traffic, it may beneficial to use not one NetGuards, but a farm of NetGuard. However, one should noticed that the problem of attacks, and special spoofing attack, in most of the cases is harder when the attack is originated from outside the network. When the attack is originated from inside network, there is full information and management of the network. Hence ingress and egress filtering can be used, for dealing with spoofing attack. In cases when the origins of the attack are known, one can more easily stop the attack, by disable the origins of the attack.

³ Unlike BGP, these routing protocols adapt very quickly to topological changes, thus correcting the forwarding tables in all the routers in hundreds of milliseconds.




Figure 2: The *victim private address* is known only to trustable parts in the networks, i.e., the interfaces of routers that connected to another router or to netGuards. The routes in the network where the *victim private address* is known is marked by dashed red lines.

4 TCP Anti-spoofing

We describe here an anti-spoofing techniques which is TCP oriented. The basic principle of the idea was presented before by Checkpoint and Cisco (TCP intercept). The innovation here is in the way it is combined with the other mechanisms and its employment on the borders of the hosting AS. This anti-spoofing mechanism authenticates the genuine of the source address of the flow, based on the SYN mechanism of TCP. Doing so each guard in effect becomes a very low level TCP proxy for the victim. When a client wishes to open a TCP connection with a server, it sends a SYN request, notifying the server about its attempt to open a new connection. The server authenticates the client source address by sending the client a random number. Then, the server waits to receive this random number back from the source. A spoofed source cannot repeat the number, and hence any connection between the server and a spoofed source is dismissed. (see figure 2).

The SYN mechanism or the connection establishment (three way handshake) is also one of the known denial of service attack methods. In this attack a huge number of spoofed SYN-requests are being sent to the server. Each such request must be buffered and kept by the server for a period of time (30 seconds by the standard) until its corresponding SYN-ACK is received. The SYN-request buffer at the server overfills which at worse brings the server down and at best causing the server to loose good genuine requests to open new connections.

Each of our guard machines is a special purpose machine that among other things performs the connection establishment on behalf of the victim. Being special purpose it can handle a huge number of connections at very high speeds (supporting OC-192

lines). Moreover, the distributed architecture of our system distributes the load among the different guard machines.

In the next two figures we show the sequence of messages during the three way handshake. The first figure, Figure 3, shows the normal sequence of messages during the three way handshake between a client whose IP address is 12.12.12.12. and a server whose IP address is 10.10.10.10. In Figure 4 the same process is performed but now the SYN request message is intercept by the guard machine which then performs the three way handshake on behalf of the server. Only after the guard machine receives the correct SYN-ACK message from the client it opens the corresponding connection with the server and starts to function as a proxy between the client and the server.

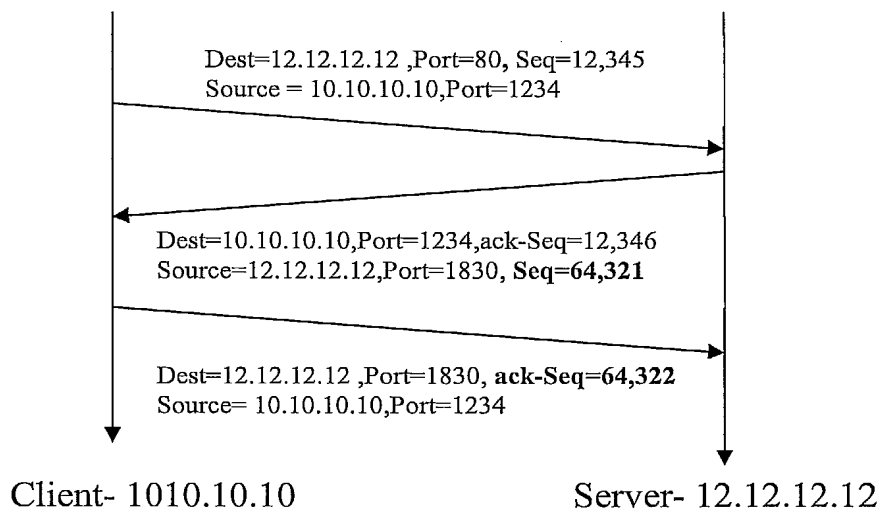
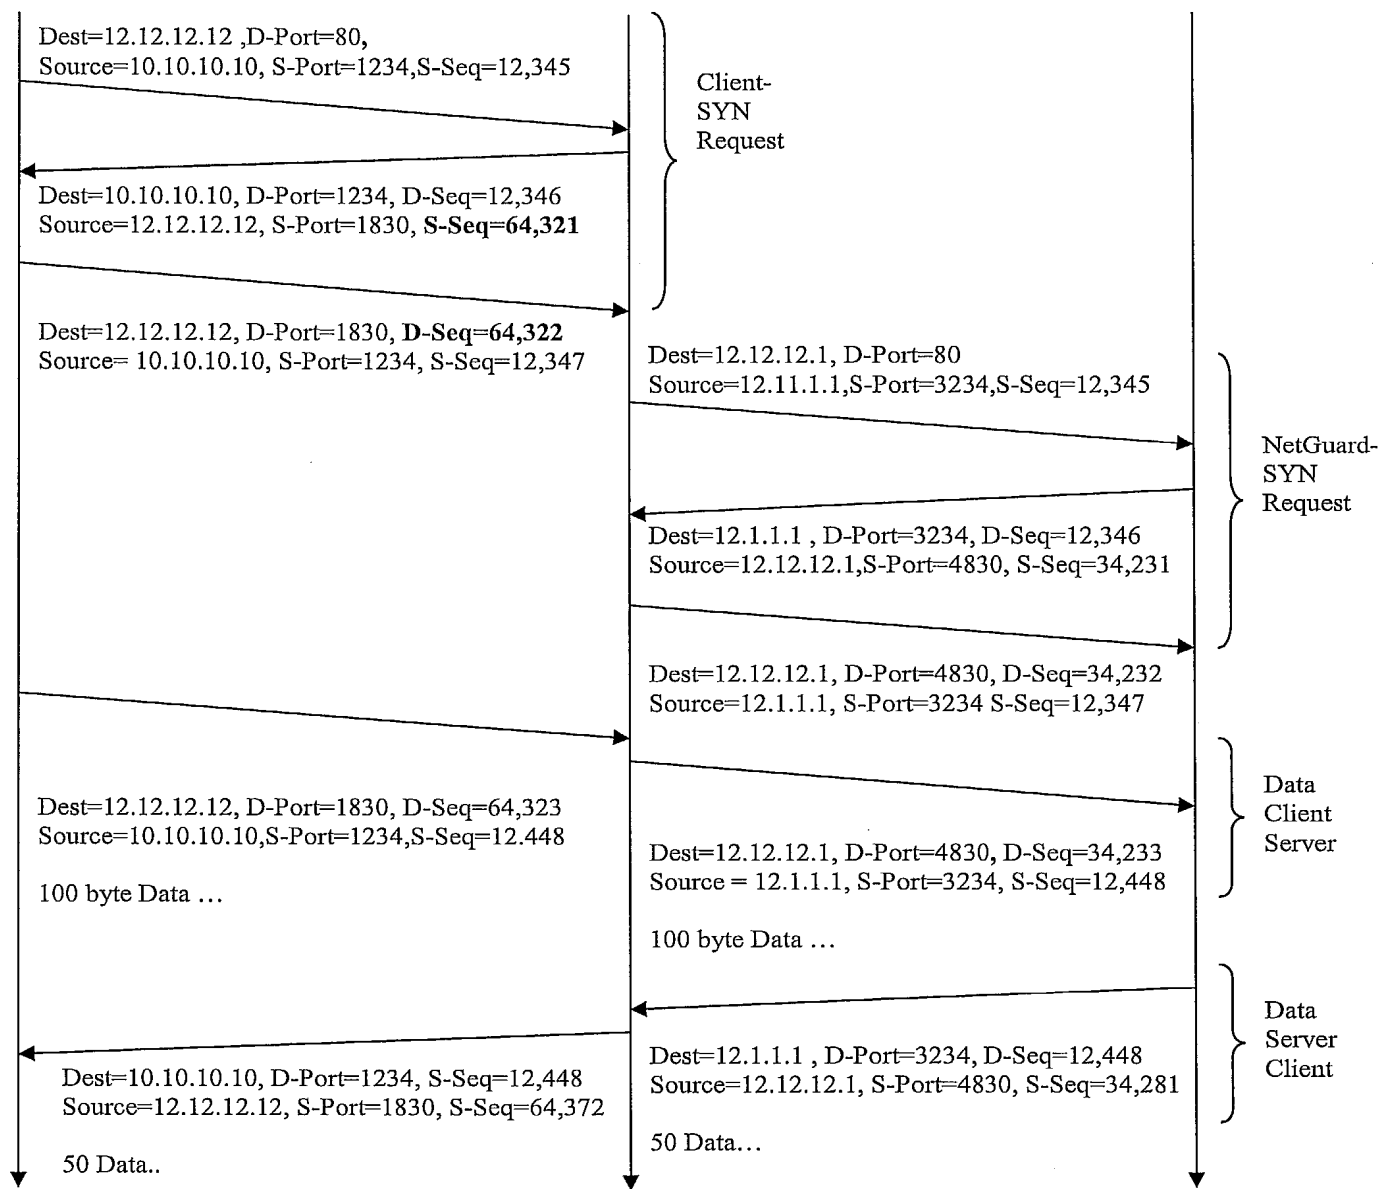


Fig 3: The SYN request.

EXHIBIT C



Client- 10.10.10.10

.1 NetGuard – 12.1.1.1
 Play the rule of 12.12.12.12
 12.12.12.12 is the server
 public address

Server- 12.12.12.1
 this address is the
 private address of the
 victim

5. **Ingress filtering**

All bla bal bla bla bla authentication and was not stopped by the

6. **Attack Identification, Recognition and Isolation via the Statistical Recognition Unit.**

All the victim traffic that has passed the anti-spoofing authentication and was not stopped by the filter flows through the statistical unit. The statistical unit analyzes the traffic and identifies malicious sources (i.e., compromised sources), and provides operational rules for blocking the attack without disturbing innocent genuine traffic. The basic principle behind the unit's operation is that the pattern of traffic originating from a black-hat daemon drastically differs from the pattern generated by such a source during normal operation. In contrast, traffic patterns of "innocent" sources during an attack resemble their traffic at normal times. This principle is used to identify the attack sources and provide guidelines for their blockage by either the filter or the access lists of the routers. For example, the volume of a traffic from an attacking daemon, the distribution of packet sizes, port numbers, the distribution of the packets inter arrival times, and the ratio of inbound and outbound traffic are all parameters that may indicate that a source (client) is an attacking daemon.

The statistical unit has two major tasks:

1. Learning the traffic patterns during normal operation, i.e., when no attack is being mounted. These patterns are used while defending a victim during an attack to compare with the actual traffic in order to distinguish the malicious traffic from genuine traffic. We consider three possible ways in which this learning can be done: (1) Using the routers NetFlow data, (2) Analyzing the server logs at the victim server, and (3) Analyzing the potential victim traffic at the guard by having the traffic diverted to the guard from time to time for randomly sampling it.
2. Monitoring the victim traffic during an attack to identify and isolate the malicious traffic from the good genuine traffic. The identity of the attacking host is then given to the filter or the neighboring routers that would then drop any packet arriving from that host.

5.1 **Network flows and traffic classification**

The basic element studied by the statistical unit is a flow. Each flow is a sequence of packets belonging to the same connection. In the most general way a flow is identified by the following parameters: Source IP address, Source port, Destination port, Protocol type, time of day and day of week of connection creation. The destination IP address is implied since we collect all the information per destination address. For each such flow the traffic volume is registered.

Keeping all of the above information is infeasible since it requires an unacceptable amount of memory. However we employ learning methods to study the basic characteristics of the traffic destined to each destination and keep these key parameters succinctly, in an efficient way. Essentially the learning method studies the typical behavior of groups of users that interact with the destination. For example, a typical web site is accessed either by individual users sitting behind a host (pc), by a group of users sharing one multi user time sharing host, or by a group of users sitting behind a proxy. For each such group its typical behavior is studied. Other types of users are possible such as web crawlers (for search engines) and monitoring servers such as keynote (www.keynote.com). Furthermore, for the largest group, i.e. the group of individual users, their identities (IP address) is not kept. Each source which is not included in the other groups is assumed to be an individual source. On the other hand, for the group of proxy machines that access the destination, the individual IP address of each is kept in a trie like data-structure. For other groups of users only their IP address may be needed since their traffic would be blocked from the beginning during an attack. Henceforth, the rest of this section considers types of users and the characteristic of flows originating from such users.

The basic parameters characterizing each user group are:

1. **Traffic volume distribution:** These include the **mean** and **variance** of the traffic such a user generates.
2. **Port numbers distribution:** Source port number distribution, and destination port number distribution.
3. **Periodicity:** Sources will be examined for the periodicity of their requests. It is likely that malicious sources act in a relatively periodic manner, while innocent sources act not in a regular manner.
4. **Packet Properties:** The distribution of packet sizes.

5.2 Learning Traffic Characteristics

There are three possible ways, that we consider, to learn and analyze the traffic characteristics of a particular target:

1. Sampling a **fraction** α of the **packets** ($0 < \alpha \leq 1$) traversing the lines on route to the target and then classifying the packets according to the flow id and time of day and day of week.
Notice that setting $\alpha=1$ requires the unit to process every packet and thus imposes high load on it while providing the best statistical measure. Lower values of α reduce the load posed on the unit while potentially somewhat degrading the statistical measure. The fraction α , therefore, will be a parameter that will be set so that enough statistical knowledge can be gained without over-loading the system.
2. Utilizing **server logs** collected by the defended target. These typically contain information about the activity being applied on the target. For example, WEB sites, which are likely to form the main body of potential targets, keep logs that record all the document requests sent to the site (including their source address, time of the day and other parameters). Processing of these logs by ZZZ yields a very accurate measure of the statistics of network flow volumes

(measured in packets per second, as in a) above). The potential draw backs of this method are first that being collected at the target it is not immediately clear which information is relevant to which guarding point, and second, the pattern seen by the target may be slightly different from the pattern seen at the network boarder. However neither is a real problem and the first one may be a feature, since network routes change and the traffic may enter the network from a different point in any event.

3. Analyzing netflow data collected from the appropriated routers. This option requires the backbone provider to enable netflow and process it with our learning applications. This method has some limitation but none seems prohibitory. The limitations are that netflow aggregates information for each flow in intervals of a few minutes (typically 5 minutes intervals), and in this intervals it does not maintain the sizes of individual packets. Rather, it counts the total number of packets and bytes passed in this interval for each flow.

5.3 **Traffic Monitoring and Analysis at Attack Time**

In attack time while the guard machine defends a victim it monitors the victim traffic, classifies its traffic (incoming and outgoing) and compares the traffic to the normal traffic in order to detect the malicious traffic. Notice that during an attack information is collected only on the current flows. The information about well behaving flows is not kept more than small number of minutes.

1. **Online traffic volume collection at attack time:** This module collects the statistics of the traffic destined to the target(s) in attack time. Notice that in this sense, its measures are similar to the measures collected in approach 1i above. The classification of the traffic, in general, is similar to that conducted in the learning phase but may be controlled/guided by external intervention. Such intervention is enacted if some additional knowledge on the attack type is gained from other sources (e.g., human-aided identification) and can be utilized by the unit.
2. **Attack Analysis:** Is conducted in attack time and is responsible to compare the statistical data learned with the current traffic volume and generate rules for traffic blockage. The output of this unit, in general, will consist of a list of items for each of which three parameters will be provided:
 - a. **Network flow**, identified by a combination of source IP address (can be prefixed), destination IP address, destination port number, protocol type (one may consider blockage that disregards port numbers, i.e., all the traffic originating from a compromised IP address, be it a proxy or a host).
 - b. **Duration**, identifying the duration for which that class will be blocked.

The analysis is based on the statistical parameters of the data and aims at keeping the target traffic at normal loads by blocking the most “suspicious” traffic streams. Blocking rules are based on maximizing

the likelihood of blocking malicious traffic while minimizing the likelihood of blocking innocent traffic.

5.3.1 Statistical Recognition of Data “Innocence”

NETGUARDS uses two major properties of network flows to identify malicious traffic: a) Traffic pattern, and b) Traffic volume. Below we describe the recognition approaches based on these factors.

5.3.1.1 Recognition of Traffic Pattern

Several aspects of traffic pattern are examined:

1) **Source “IP geography” proximity:** Sources will be classified into classes that resemble the “IP geography”, that is IP addresses that reside on neighboring networks (using IP address prefix) are classified in the same class. A class that generates a relatively large volume of requests is suspected as being malicious. Notice, that such “malicious classes” are likely to form if the attacker planted a collection of daemons in the same network, and this network does not use a proxy.

2) **Periodicity:** Sources will be examined for the periodicity of their requests. It is likely that malicious sources act in a relatively periodic manner, while innocent sources act in more irregular pattern.

3) **Packet Properties:** Sources will be examined for repetitive properties of their packets. For example the distribution of packet size. It is likely that malicious sources generate packets of identical properties (e.g. – all packets of same size) while innocent sources will generate packets of more random nature. Other properties include port number distributions.

5.3.1.2 Recognition of Traffic Volume

Traffic volume recognition is used to identify malicious sources that transmit **large volumes** of data which **significantly differ** from their normal volume. Specifically, we classify Internet data sources to *small sources* and *large sources*. The former relates to individual IP addresses whose traffic volume is normally tiny. The latter relates to Proxy traffic or Spider (Crawler) traffic⁴ whose volume is drastically higher.

ZZZ keeps individual volume parameters for each of the large sources. Individual parameters are not kept for the small sources; rather a single fixed small number (related to their mean volume averaged over all these sources) will be recorded. At attack time the traffic volumes of individual flows will be measured and compared to their recorded volume. Flows whose volume drastically differs (upwards) from their recorded measure are marked as being malicious.

⁴ The traffic volume resulting from a Spider access is normally higher than that of a single human user, especially on large WEB sites. The reason is that a spider scans the whole site, leading to hundred or thousands of requests while a human client requests tens of pages or less, on average.

The mathematical formulation of this procedure is as follows: Given are K classes of flows, indexed $1, 2, \dots, K$, and characterized by the mean (μ_i) and the variance (σ_i) of their learned volume, and by their current volume (X_i). We would like to identify the classes with the largest deviation from their corresponding expected volume. Let $Y_i = (X_i - \mu_i) / \sigma_i$. We will sort the classes by the value of Y_i and recommend blocking the classes with the largest values of Y_i .

5.3.2 Time accumulating traffic volume recognition and “controlled” denial of service

It is important to recognize that the effectiveness of volume recognition increases with the time duration along which it is implemented. This is correct since the variance of total data volume generated by a source during a period of duration T decreases in T . For example, it is expected that the average amount of traffic generated by a small source during a period of 1 hour will be *very small*. However, at certain epochs, it is expected that the average amount of traffic generated by the same source during a period of 1 minute can be rather large. (up to 60 times larger than that of the 1 hour average).

For this reason ZZZ will implement the following unique recognition and traffic screening mechanism. For source i , let $S_i(t)$ denote the amount traffic generated by the source during the interval $(0, t)$ (where we assume that the attack starts at time 0). We then set at time t : $X_i(t) = S_i(t) / t$ and apply the above screening mechanism.

This mechanism has the following properties:

1. For a small value of t (that is, at the first few minutes of the attack) a sophisticated attacker might cause significant number of innocent users denial of service. This is due to the fact that the attacker may inflict a load that resembles that of an innocent client, and thus the attacker is not distinguishable from the innocent client. At this stage, ZZZ may block some innocent clients and some attackers. Using this action, for a short period of time, some innocent clients may be denied of service but ZZZ protects the site from going down.
2. As t increases more and more malicious sources are identified and blocked and fewer innocent sources are blocked. This is since the malicious sources have posed large amount of accumulated load. Thus as time progresses less and less innocent clients are denied service. In fact, after relatively short period all malicious sources will be denied service while the

EXHIBIT C

innocent sources will receive full regular service.

Example: Consider the traffic volume generated on the web site of the Nagano server (Feb 98). It had 11,665,713 requests made over a period of 24 hours by 59,582 clients. Assuming uniform distribution of clients over the day, this implies about 2500 clients per hour and 500 clients per 12-minute interval. An attacker who uses 500 sophisticated daemons (which imitate a normal client) will look innocent at the first 12 minutes interval. At this period ZZZ will block 50% of the innocent clients and 50% of the daemons. However, after 24 minutes the daemons will generate significantly more traffic than an innocent client and thus almost all of the traffic blocked will be that of malicious daemons.

5. **Attack Identification, Recognition and Isolation via the Statistical Recognition Unit.**

All the victim traffic that has passed the anti-spoofing authentication and was not stopped by the filter flows through the statistical unit. The statistical unit analyzes the traffic and identifies malicious sources (i.e., compromised sources), and provides operational rules for blocking the attack without disturbing innocent genuine traffic. The basic principle behind the unit's operation is that the pattern of traffic originating from a black-hat daemon drastically differs from the pattern generated by such a source during normal operation. In contrast, traffic patterns of "innocent" sources during an attack resemble their traffic at normal times. This principle is used to identify the attack sources and provide guidelines for their blockage by either the filter or the access lists of the routers. For example, the volume of a traffic from an attacking daemon, the distribution of packet sizes, port numbers, the distribution of the packets inter arrival times, and the ratio of inbound and outbound traffic are all parameters that may indicate that a source (client) is an attacking daemon.

The statistical unit has two major tasks:

1. Learning the traffic patterns during normal operation, i.e., when no attack is being mounted. These patterns are used while defending a victim during an attack to compare with the actual traffic in order to distinguish the malicious traffic from genuine traffic. We consider three possible ways in which this learning can be done: (1) Using the routers NetFlow data, (2) Analyzing the server logs at the victim server, and (3) Analyzing the potential victim traffic at the guard by having the traffic diverted to the guard from time to time for randomly sampling it.
2. Monitoring the victim traffic during an attack to identify and isolate the malicious traffic from the good genuine traffic. The identity of the attacking host is then given to the filter or the neighboring routers that would then drop any packet arriving from that host.

5.1 **Network flows and traffic classification**

The basic element studied by the statistical unit is a flow. Each flow is a sequence of packets belonging to the same connection. In the most general way a flow is identified by the following parameters: Source IP address, Source port, Destination port, Protocol type, time of day and day of week of connection creation. The destination IP address is implied since we collect all the information per destination address. For each such flow the traffic volume is registered.

Keeping all of the above information is infeasible since it requires an unacceptable amount of memory. However we employ learning methods to study the basic characteristics of the traffic destined to each destination and keep these key parameters succinctly, in an efficient way. Essentially the learning method studies the typical behavior of groups of users that interact with the destination. For example, a typical web site is accessed either by individual users sitting behind a host (pc), by a group of users sharing one multi user time sharing host, or by a group of users sitting behind a proxy. For each such group its typical behavior is studied. Other types of

users are possible such as web crawlers (for search engines) and monitoring servers such as keynote (www.keynote.com). Furthermore, for the largest group, i.e. the group of individual users, their identities (IP address) is not kept. Each source which is not included in the other groups is assumed to be an individual source. On the other hand, for the group of proxy machines that access the destination, the individual IP address of each is kept in a trie like data-structure. For other groups of users only their IP address may be needed since their traffic would be blocked from the beginning during an attack. Henceforth, the rest of this section considers types of users and the characteristic of flows originating from such users.

The basic parameters characterizing each user group are:

1. **Traffic volume distribution:** These include the **mean** and **variance** of the traffic such a user generates.
2. **Port numbers distribution:** Source port number distribution, and destination port number distribution.
3. **Periodicity:** Sources will be examined for the periodicity of their requests. It is likely that malicious sources act in a relatively periodic manner, while innocent sources act not in a regular manner.
4. **Packet Properties:** The distribution of packet sizes.

5.2 Learning Traffic Characteristics

There are three possible ways, that we consider, to learn and analyze the traffic characteristics of a particular target:

1. Sampling a **fraction** α of the **packets** ($0 < \alpha \leq 1$) traversing the lines on route to the target and then classifying the packets according to the flow id and time of day and day of week.
 Notice that setting $\alpha=1$ requires the unit to process every packet and thus imposes high load on it while providing the best statistical measure. Lower values of α reduce the load posed on the unit while potentially somewhat degrading the statistical measure. The fraction α , therefore, will be a parameter that will be set so that enough statistical knowledge can be gained without over-loading the system.
2. Utilizing **server logs** collected by the defended target. These typically contain information about the activity being applied on the target. For example, WEB sites, which are likely to form the main body of potential targets, keep logs that record all the document requests sent to the site (including their source address, time of the day and other parameters). Processing of these logs by ZZZ yields a very accurate measure of the statistics of network flow volumes (measured in packets per second, as in a) above). The potential draw backs of this method are first that being collected at the target it is not immediately clear which information is relevant to which guarding point, and second, the pattern seen by the target may be slightly different from the pattern seen at the network boarder. However neither is a real problem and the first one may be a feature, since network routes change and the traffic may enter the network from a different point in any event.

3. Analyzing netflow data collected from the appropriated routers. This option requires the backbone provider to enable netflow and process it with our learning applications. This method has some limitation but none seems prohibitory. The limitations are that netflow aggregates information for each flow in intervals of a few minutes (typically 5 minutes intervals), and in this intervals it does not maintain the sizes of individual packets. Rather, it counts the total number of packets and bytes passed in this interval for each flow.

5.3 *Traffic Monitoring and Analysis at Attack Time*

In attack time while the guard machine defends a victim it monitors the victim traffic, classifies its traffic (incoming and outgoing) and compares the traffic to the normal traffic in order to detect the malicious traffic. Notice that during an attack information is collected only on the current flows. The information about well behaving flows is not kept more than small number of minutes.

1. **Online traffic volume collection at attack time:** This module collects the statistics of the traffic destined to the target(s) in attack time. Notice that in this sense, its measures are similar to the measures collected in approach 1i above. The classification of the traffic, in general, is similar to that conducted in the learning phase but may be controlled/guided by external intervention. Such intervention is enacted if some additional knowledge on the attack type is gained from other sources (e.g., human-aided identification) and can be utilized by the unit.
2. **Attack Analysis:** Is conducted in attack time and is responsible to compare the statistical data learned with the current traffic volume and generate rules for traffic blockage. The output of this unit, in general, will consist of a list of items for each of which three parameters will be provided:
 - a. **Network flow**, identified by a combination of source IP address (can be prefixed), destination IP address, destination port number, protocol type (one may consider blockage that disregards port numbers, i.e., all the traffic originating from a compromised IP address, be it a proxy or a host).
 - b. **Duration**, identifying the duration for which that class will be blocked.

The analysis is based on the statistical parameters of the data and aims at keeping the target traffic at normal loads by blocking the most “suspicious” traffic streams. Blocking rules are based on maximizing the likelihood of blocking malicious traffic while minimizing the likelihood of blocking innocent traffic.

5.3.1 Statistical Recognition of Data “Innocence”

NETGUARDS uses two major properties of network flows to identify malicious traffic: a) Traffic pattern, and b) Traffic volume. Below we describe the recognition approaches based on these factors.

5.3.1.1

Recognition of Traffic Pattern

Several aspects of traffic pattern are examined:

1) **Source “IP geography” proximity:** Sources will be classified into classes that resemble the “IP geography”, that is IP addresses that reside on neighboring networks (using IP address prefix) are classified in the same class. A class that generates a relatively large volume of requests is suspected as being malicious. Notice, that such “malicious classes” are likely to form if the attacker planted a collection of daemons in the same network, and this network does not use a proxy.

2) **Periodicity:** Sources will be examined for the periodicity of their requests. It is likely that malicious sources act in a relatively periodic manner, while innocent sources act in more irregular pattern.

3) **Packet Properties:** Sources will be examined for repetitive properties of their packets. For example the distribution of packet size. It is likely that malicious sources generate packets of identical properties (e.g. – all packets of same size) while innocent sources will generate packets of more random nature. Other properties include port number distributions.

5.3.1.2

Recognition of Traffic Volume

Traffic volume recognition is used to identify malicious sources that transmit **large volumes** of data which **significantly differ** from their normal volume. Specifically, we classify Internet data sources to *small sources* and *large sources*. The former relates to individual IP addresses whose traffic volume is normally tiny. The latter relates to Proxy traffic or Spider (Crawler) traffic¹ whose volume is drastically higher.

ZZZ keeps individual volume parameters for each of the large sources. Individual parameters are not kept for the small sources; rather a single fixed small number (related to their mean volume averaged over all these sources) will be recorded. At attack time the traffic volumes of individual flows will be measured and compared to their recorded volume. Flows whose volume drastically differs (upwards) from their recorded measure are marked as being malicious.

The mathematical formulation of this procedure is as follows: Given are K classes of flows, indexed $1, 2, \dots, K$, and characterized by the mean (μ_i) and the variance (σ_i) of their learned volume, and by their current volume (X_i). We would like to identify the classes with the largest deviation from their corresponding expected volume. Let $Y_i = (X_i - \mu_i) / \sigma_i$. We will sort the classes

¹ The traffic volume resulting from a Spider access is normally higher than that of a single human user, especially on large WEB sites. The reason is that a spider scans the whole site, leading to hundred or thousands of requests while a human client requests tens of pages or less, on average.

by the value of Y_i and recommend blocking the classes with the largest values of Y_i .

5.3.2 Time accumulating traffic volume recognition and “controlled” denial of service

It is important to recognize that the effectiveness of volume recognition increases with the time duration along which it is implemented. This is correct since the variance of total data volume generated by a source during a period of duration T decreases in T . For example, it is expected that the average amount of traffic generated by a small source during a period of 1 hour will be *very small*. However, at certain epochs, it is expected that the average amount of traffic generated by the same source during a period of 1 minute can be rather large. (up to 60 times larger than that of the 1 hour average).

For this reason ZZZ will implement the following unique recognition and traffic screening mechanism. For source i , let $S_i(t)$ denote the amount traffic generated by the source during the interval $(0, t)$ (where we assume that the attack starts at time 0). We then set at time t : $X_i(t) = S_i(t) / t$ and apply the above screening mechanism.

This mechanism has the following properties:

1. For a small value of t (that is, at the first few minutes of the attack) a sophisticated attacker might cause significant number of innocent users denial of service. This is due to the fact that the attacker may inflict a load that resembles that of an innocent client, and thus the attacker is not distinguishable from the innocent client. At this stage, ZZZ may block some innocent clients and some attackers. Using this action, for a short period of time, some innocent clients may be denied of service but ZZZ protects the site from going down.
2. As t increases more and more malicious sources are identified and blocked and fewer innocent sources are blocked. This is since the malicious sources have posed large amount of accumulated load. Thus as time progresses less and less innocent clients are denied service. In fact, after relatively short period all malicious sources will be denied service while the innocent sources will receive full regular service.

Example: Consider the traffic volume generated on the web site of the Nagano server (Feb 98). It had 11,665,713 requests made over a period of 24 hours by 59,582 clients. Assuming uniform distribution of clients over the day,

EXHIBIT D

this implies about 2500 clients per hour and 500 clients per 12-minute interval. An attacker who uses 500 sophisticated daemons (which imitate a normal client) will look innocent at the first 12 minutes interval. At this period ZZZ will block 50% of the innocent clients and 50% of the daemons. However, after 24 minutes the daemons will generate significantly more traffic than an innocent client and thus almost all of the traffic blocked will be that of malicious daemons.

5. Attack Identification, Recognition and Isolation via the Statistical Recognition Unit.

All the victim traffic that has passed the anti-spoofing authentication and was not stopped by the filter flows through the statistical unit. The statistical unit analyzes the traffic and identifies malicious sources (i.e., compromised sources), and provides operational rules for blocking the attack without disturbing innocent genuine traffic. The basic principle behind the unit's operation is that the pattern of traffic originating from a black-hat daemon drastically differs from the pattern generated by such a source during normal operation. In contrast, traffic patterns of "innocent" sources during an attack resemble their traffic at normal times. This principle is used to identify the attack sources and provide guidelines for their blockage by either the filter or the access lists of the routers. For example, the volume of a traffic from an attacking daemon, the distribution of packet sizes, port numbers, the distribution of the packets inter arrival times, and the ratio of inbound and outbound traffic are all parameters that may indicate that a source (client) is an attacking daemon.

The statistical unit has two major tasks:

1. Learning the traffic patterns during normal operation, i.e., when no attack is being mounted. These patterns are used while defending a victim during an attack to compare with the actual traffic in order to distinguish the malicious traffic from genuine traffic. We consider three possible ways in which this learning can be done: (1) Using the routers NetFlow data, (2) Analyzing the server logs at the victim server, and (3) Analyzing the potential victim traffic at the guard by having the traffic diverted to the guard from time to time for randomly sampling it.
2. Monitoring the victim traffic during an attack to identify and isolate the malicious traffic from the good genuine traffic. The identity of the attacking host is then given to the filter or the neighboring routers that would then drop any packet arriving from that host.

5.1 Network flows and traffic classification

The basic element studied by the statistical unit is a flow. Each flow is a sequence of packets belonging to the same connection. In the most general way a flow is identified by the following parameters: Source IP address, Source port, Destination port, Protocol type, time of day and day of week of connection creation. The destination IP address is implied since we collect all the information per destination address. For each such flow the traffic volume is registered.

Keeping all of the above information is infeasible since it requires an unacceptable amount of memory. However we employ learning methods to study the basic characteristics of the traffic destined to each destination and keep these key parameters succinctly, in an efficient way. Essentially the learning method studies the typical behavior of groups of users that interact with the destination. For example, a typical web site is accessed either by individual users sitting behind a host (pc), by a group of users sharing one multi user time sharing host, or by a group of users sitting behind a proxy. For each such group its typical behavior is studied. Other types of

users are possible such as web crawlers (for search engines) and monitoring servers such as keynote (www.keynote.com). For the individual users, their identities (IP address) are not kept to save storage. Each source which is not included in the other groups is assumed to be an individual source. On the other hand, for the group of proxy machines that access the destination, the individual IP address of each is kept in a trie like data-structure. For other groups of users only their IP address may be needed since their traffic would be blocked from the beginning during an attack. Henceforth, the rest of this section considers types of users and the characteristic of flows originating from such users.

The basic parameters characterizing each user group are:

1. **Traffic volume distribution:** These include the **mean** and **variance** of the traffic such a user generates.
2. **Port numbers distribution:** Source port number distribution, and destination port number distribution.
3. **Periodicity:** Sources will be examined for the periodicity of their requests. It is likely that malicious sources act in a relatively periodic manner, while innocent sources act not in a regular manner.
4. **Packet Properties:** The distribution of packet sizes.

5.2 Learning Traffic Characteristics

There are three possible ways, that we consider, to learn and analyze the traffic characteristics of a particular target:

1. Sampling a **fraction** α of the **packets** ($0 < \alpha \leq 1$) traversing the lines on route to the target and then classifying the packets according to the flow id and time of day and day of week.
Notice that setting $\alpha=1$ requires the unit to process every packet and thus imposes high load on it while providing the best statistical measure. Lower values of α reduce the load posed on the unit while potentially somewhat degrading the statistical measure. The fraction α , therefore, will be a parameter that will be set so that enough statistical knowledge can be gained without over-loading the system.
2. Utilizing **server logs** collected by the defended target. These typically contain information about the activity being applied on the target. For example, WEB sites, which are likely to form the main body of potential targets, keep logs that record all the document requests sent to the site (including their source address, time of the day and other parameters). Processing of these logs by ZZZ yields a very accurate measure of the statistics of network flow volumes (measured in packets per second, as in a) above). The potential draw backs of this method are first that being collected at the target it is not immediately clear which information is relevant to which guarding point, and second, the pattern seen by the target may be slightly different from the pattern seen at the network border. However neither is a real problem and the first one may be a feature, since network routes change and the traffic may enter the network from a different point in any event.

3. Analyzing netflow data collected from the appropriated routers. This option requires the backbone provider to enable netflow and process it with our learning applications. This method has some limitation but none seems prohibitory. The limitations are that netflow aggregates information for each flow in intervals of a few minutes (typically 5 minutes intervals), and in this intervals it does not maintain the sizes of individual packets. Rather, it counts the total number of packets and bytes passed in this interval for each flow.

5.3 Traffic Monitoring and Analysis at Attack Time

In attack time while the guard machine defends a victim it monitors the victim traffic, classifies its traffic (incoming and outgoing) and compares the traffic to the normal traffic in order to detect the malicious traffic. Notice that during an attack information is collected only on the current flows. The information about well behaving flows is not kept more than small number of minutes.

1. **Online traffic volume collection at attack time:** This module collects the statistics of the traffic destined to the target(s) in attack time. Notice that in this sense, its measures are similar to the measures collected in approach 1i above. The classification of the traffic, in general, is similar to that conducted in the learning phase but may be controlled/guided by external intervention. Such intervention is enacted if some additional knowledge on the attack type is gained from other sources (e.g., human-aided identification) and can be utilized by the unit.
2. **Attack Analysis:** Is conducted in attack time and is responsible to compare the statistical data learned with the current traffic volume and generate rules for traffic blockage. The output of this unit, in general, will consist of a list of items for each of which three parameters will be provided:
 - a. **Network flow**, identified by a combination of source IP address (can be prefixed), destination IP address, destination port number, protocol type (one may consider blockage that disregards port numbers, i.e., all the traffic originating from a compromised IP address, be it a proxy or a host).
 - b. **Duration**, identifying the duration for which that class will be blocked.

The analysis is based on the statistical parameters of the data and aims at keeping the target traffic at normal loads by blocking the most “suspicious” and “harmful” traffic streams. Blocking rules are based on maximizing the likelihood of blocking malicious traffic while minimizing the likelihood of blocking innocent traffic.

5.3.1 Statistical Recognition of Data “Innocence”

NETGUARDS uses two major properties of network flows to identify malicious traffic: a) Traffic pattern, and b) Traffic volume. Below we describe the recognition approaches based on these factors.

5.3.1.1 Recognition of Traffic Pattern

Several aspects of traffic pattern are examined:

1) **Source “IP geography” proximity:** Sources will be classified into classes that resemble the “IP geography”, that is IP addresses that reside on neighboring networks (using IP address prefix) are classified in the same class. A class that generates a relatively large volume of requests is suspected as being malicious. Notice, that such “malicious classes” are likely to form if the attacker planted a collection of daemons in the same network, and this network does not use a proxy.

2) **Periodicity:** Sources will be examined for the periodicity of their requests. It is likely that malicious sources act in a relatively periodic manner, while innocent sources act in more irregular pattern.

3) **Packet Properties:** Sources will be examined for repetitive properties of their packets. For example the distribution of packet size. It is likely that malicious sources generate packets of identical properties (e.g. – all packets of same size) while innocent sources will generate packets of more random nature. Other properties include port number distributions.

5.3.1.2 Recognition of Traffic Volume

Traffic volume recognition is used to identify malicious sources that transmit **large volumes** of data which **significantly differ** from their normal volume. Specifically, we classify Internet data sources to *small sources* and *large sources*. The former relates to individual IP addresses whose traffic volume is normally tiny. The latter relates to Proxy traffic or Spider (Crawler) traffic¹ whose volume is drastically higher.

ZZZ keeps individual volume parameters for each of the large sources. Individual parameters are not kept for the small sources; rather a single fixed small number (related to their mean volume averaged over all these sources) will be recorded. At attack time the traffic volumes of individual flows will be measured and compared to their recorded volume. Flows whose volume drastically differs (upwards) from their recorded measure are marked as being malicious.

The mathematical formulation of this procedure is as follows: Given are K classes of flows, indexed $1, 2, \dots, K$, and characterized by the mean (μ_i) and the variance (σ_i) of their learned volume, and by their current volume (X_i). We would like to identify the classes with the largest deviation from their corresponding expected volume. Let $Y_i = (X_i - \mu_i) / \sigma_i$. We will sort the classes by the value of Y_i and recommend blocking the classes with the largest values of Y_i . Blockage of sources will be done sequentially until the total volume of sources fits the traffic volume to be sustained by the network/defended site.

¹ The traffic volume resulting from a Spider access is normally higher than that of a single human user, especially on large WEB sites. The reason is that a spider scans the whole site, leading to hundred or thousands of requests while a human client requests tens of pages or less, on average.

5.3.2 Time accumulating traffic volume recognition and “controlled” denial of service

It is important to recognize that the effectiveness of volume recognition increases with the time duration along which it is implemented. This is correct since the variance of total data volume generated by a source during a period of duration T decreases in T . For example, it is expected that the average amount of traffic generated by a small source during a period of 1 hour will be *very small*. However, at certain epochs, it is expected that the average amount of traffic generated by the same source during a period of 1 minute can be rather large. (up to 60 times larger than that of the 1 hour average).

For this reason ZZZ will implement the following unique recognition and traffic screening mechanism. For source i , let $S_i(t)$ denote the amount traffic generated by the source during the interval $(0, t)$ (where we assume that the attack starts at time 0). We then set at time t : $X_i(t) = S_i(t) / t$ and apply the above screening mechanism.


This mechanism has the following properties:

1. For a small value of t (that is, at the first few minutes of the attack) a sophisticated attacker might cause significant number of innocent users denial of service. This is due to the fact that the attacker may inflict a load that resembles that of an innocent client, and thus the attacker is not distinguishable from the innocent client. At this stage, ZZZ may block some innocent clients and some attackers. Using this action, for a short period of time, some innocent clients may be denied of service but ZZZ protects the site from going down.
2. As t increases more and more malicious sources are identified and blocked and fewer innocent sources are blocked. This is since the malicious sources have posed large amount of accumulated load. Thus as time progresses less and less innocent clients are denied service. In fact, after relatively short period all malicious sources will be denied service while the innocent sources will receive full regular service.

Example: Consider the traffic volume generated on the web site of the Nagano server (Feb 98). It had 11,665,713 requests made over a period of 24 hours by 59,582 clients. Assuming uniform distribution of clients over the day, this implies about 2500 clients per hour and 500 clients per 12-minute interval. An attacker who uses 500 sophisticated daemons (which imitate a normal client) will look innocent at the first 12 minutes interval. At this period ZZZ will block 50% of the innocent clients and 50% of the daemons. However, after 24 minutes the daemons will generate significantly more traffic than an innocent client and thus almost all of the traffic blocked will be that of malicious daemons.

EXHIBIT F

Distributed Network Defense Systems



Strategic Overview

September 24, 2000


Net-guards

What is Distributed Denial of Service Attack (DDoS)? (*)

- Swamp victim & deny service from customers:
 - Exhaust resources
 - Jam routers, networks & discs....
- Distributed Attacks:
 - Many sources (daemons) bombardment
 - Master machines coordinate the attackers
- Several different versions in use: Trinoo, TFN, TFN2K and stacheldraht

Distributed Denial of Service (*)

EXHIBIT F



What is being done today?

EXHIBIT F

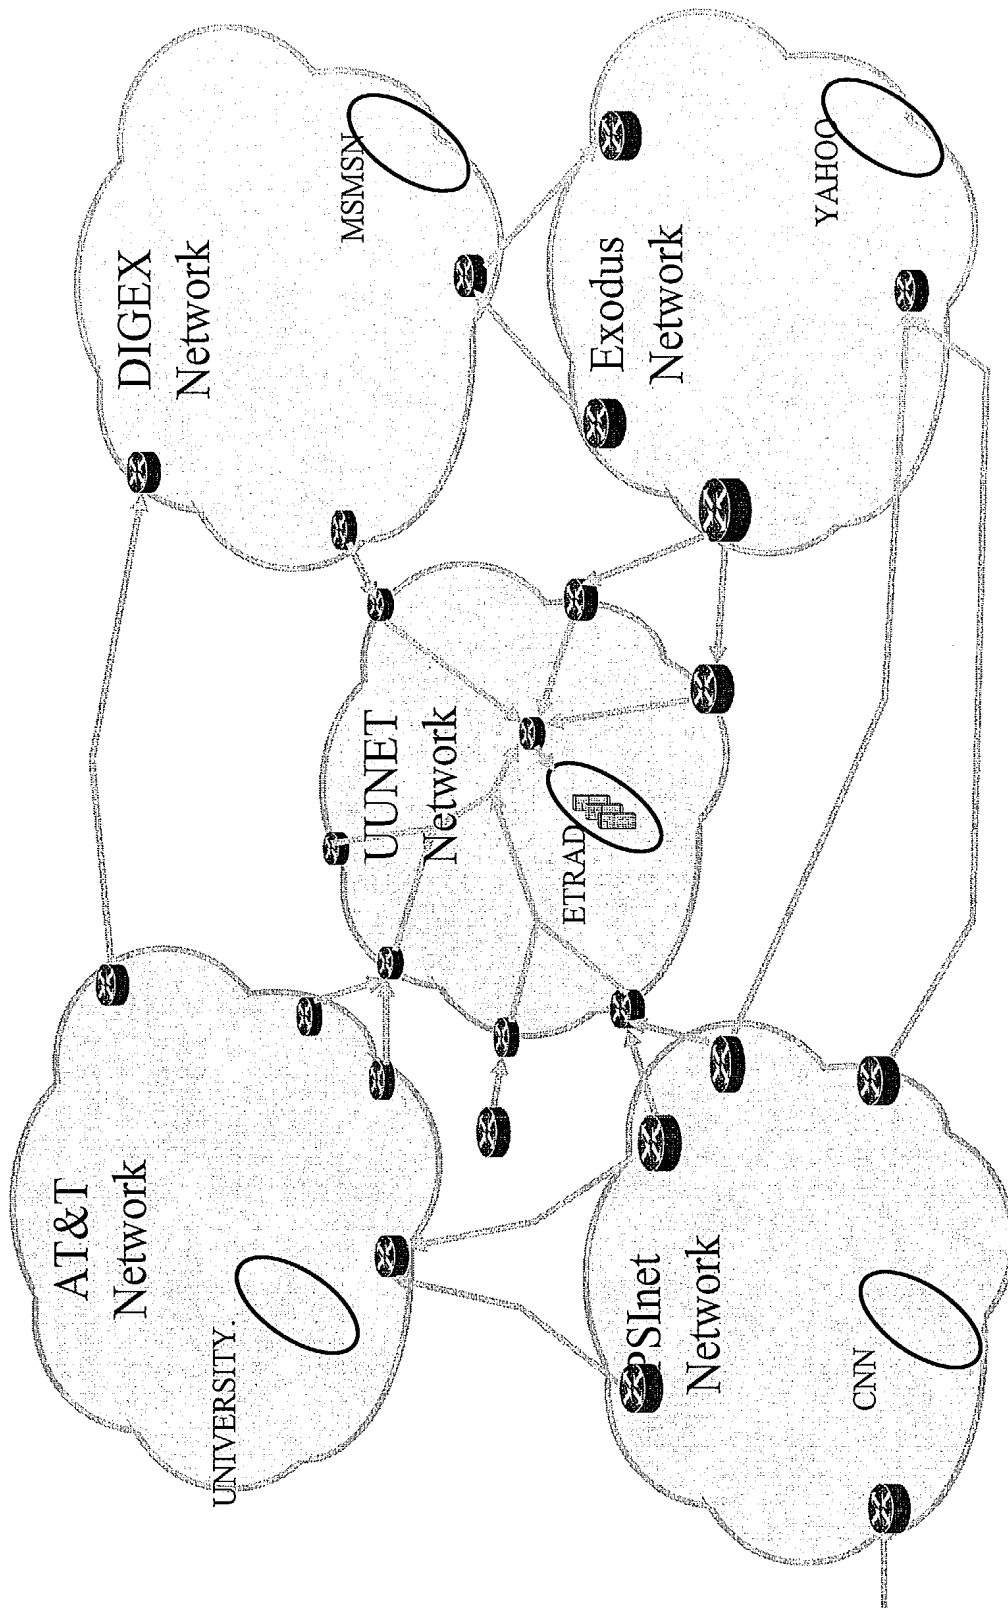
- Local on site:
 - Firewalls, IDSs, Router access lists, etc.
 - HOT SPOTS, congestion point
- Global network wide:
 - Ingress/Egress filtering and anti-spoofing
 - Anti virus scans
 - Turn off directed broadcasts
 - Enforcement problem \ Does not protect self

The Difficulties (*)


- Performance: Huge Capacity & Proc. Needs
- Defense resource not function of site size
- Solution cannot be in the site level
- "We have no real defense". Steve Bellovin

ATT Feb '00, David Dittrich WU, Aug '00

Internet structure



Solution is in the AS level (ATT, exodus, digex, MCI, ...).



Disinfectant is a cure for the disease

Solution Overview


- Detect & Alert system in each site
- Alert invokes guards sitting around
- Divert victim traffic to guards
- Guards:
 - Anti spoofing module
 - Statistical module
 - Filtering module
 - Ingress filtering management
 - Distributed termination detection module

Core Technologies

- New concept for distributed defense
- Anti spoofing & Filtering
- Diversion method
- Experience in Protocols & Distributed Algs
- New programmable Network processors
(Intel IXP1200, Ezchip, Motorola, MMC, etc.)
- Provisional patent application


Marketing Servers Models

● Dist. Guard company



Guard ●

● Gatekeeper farms



Guard Architecture

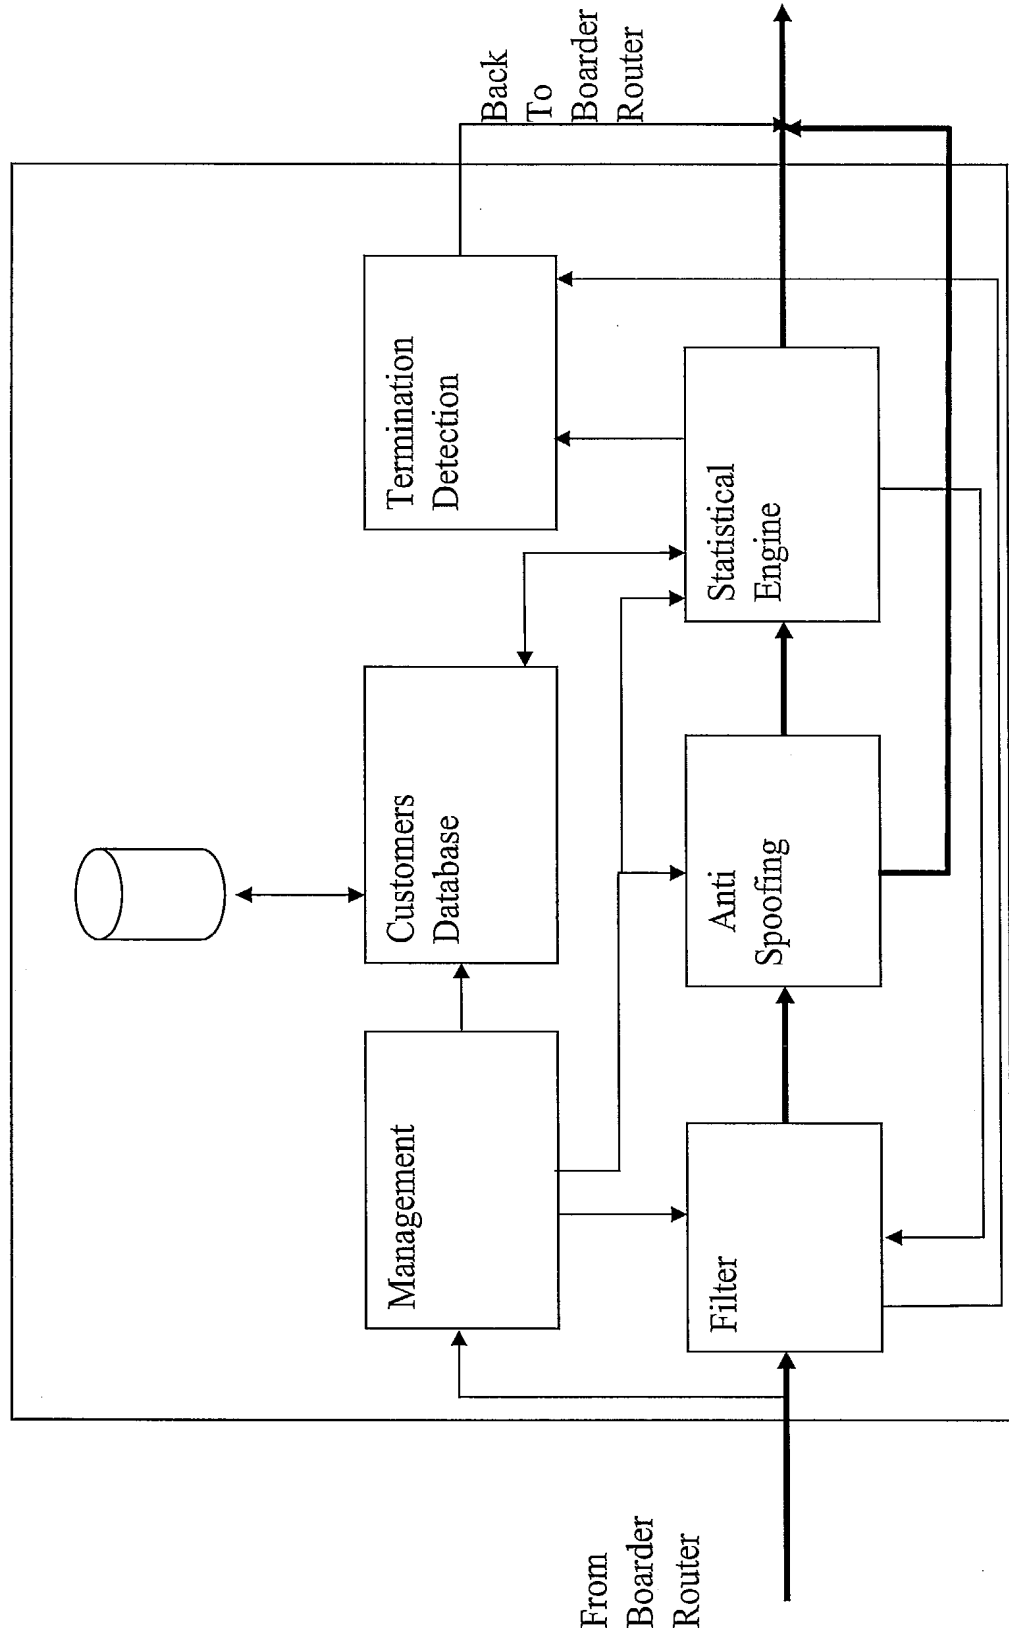


EXHIBIT G

WinZip - wawa.zip

File Actions View Jobs Options Help

New Open Favorites Add Extract Encrypt View CheckOut Wizard View Style

Name	Type	Modified	Size	Ratio	Packed	Path
0-1005-200-270...	HTML Docum...	09/09/2000 01:21	40,970	75%	10,374	wawa\
MITMultops.ps*	PostScript	09/09/2000 01:30	843,791	80%	166,420	wawa\
draft-ietf-grip-pr...	Text Document	09/09/2000 02:01	13,844	63%	5,107	wawa\
Estimates of nu...	Microsoft W...	09/09/2000 23:07	35,840	83%	6,053	wawa\ip\
hanochEsti.txt*	Text Document	09/09/2000 23:07	1,472	65%	520	wawa\ip\
dudaAug30.ppt*	Microsoft Po...	10/09/2000 01:29	253,440	69%	79,252	wawa\present\old\
sigcomm2000-3-...	PostScript	12/09/2000 01:54	348,082	3%	336,236	wawa\
EZchips.html*	HTML Docum...	13/09/2000 03:04	36,417	72%	10,319	wawa\docs\
~WRL0552.tmp*	TMP File	13/09/2000 23:44	156,160	80%	31,840	wawa\docs\
Netxx.doc*	Microsoft W...	14/09/2000 23:37	95,336	89%	10,474	wawa\ip\
DDoSprop.html*	HTML Docum...	15/09/2000 00:04	12,690	66%	4,258	wawa\
Copy of dada2....	Microsoft Po...	15/09/2000 05:33	267,264	69%	81,701	wawa\present\old\
net100.pdf*	Adobe Acro...	16/09/2000 00:08	306,786	36%	194,814	wawa\docs\
NS-1000_ds.pdf*	Adobe Acro...	16/09/2000 00:17	609,856	8%	561,671	wawa\docs\
ns1000_wpapaper...	Adobe Acro...	16/09/2000 00:31	70,604	26%	52,262	wawa\docs\
mafiaboy.asp*	ASP File	16/09/2000 13:08	46,261	74%	12,149	wawa\
Copy (2) of dad...	Microsoft Po...	17/09/2000 05:07	264,192	69%	81,353	wawa\present\old\
BP_v1Sep_00.d...	Microsoft W...	17/09/2000 05:16	76,288	81%	14,847	wawa\bp\
Statistical-paten...	Microsoft W...	18/09/2000 00:00	84,992	78%	18,441	wawa\ip\
Statistical-paten...	Microsoft W...	18/09/2000 21:20	55,296	73%	14,656	wawa\ip\
~WRL1508.tmp*	TMP File	19/09/2000 01:19	440,832	79%	93,183	wawa\docs\
bluetooth.ppt*	Microsoft Po...	19/09/2000 08:27	1,391,104	59%	572,768	wawa\
PcMgFirewalls.ht...	HTML Docum...	19/09/2000 20:50	42,336	80%	8,634	wawa\
ZDFrWlsddos.ht...	HTML Docum...	19/09/2000 20:52	47,813	78%	10,433	wawa\
ZDSpoofProof.h...	HTML Docum...	19/09/2000 20:53	47,106	77%	11,055	wawa\
denial_services...	HTML Docum...	19/09/2000 21:19	11,701	71%	3,416	wawa\
~WRL3100.tmp*	TMP File	19/09/2000 22:09	196,096	81%	37,275	wawa\docs\
AstaArtcl.html*	HTML Docum...	20/09/2000 07:13	34,976	74%	9,114	wawa\
exec2-sent.doc*	Microsoft W...	21/09/2000 21:17	60,416	78%	13,191	wawa\bp\
~WRL2764.tmp*	TMP File	22/09/2000 21:13	1,171,968	77%	268,408	wawa\docs\
RvrsPxy53-2.ps*	PostScript	24/09/2000 13:25	49,369	55%	22,231	wawa\
Copy (3) of dad...	Microsoft Po...	24/09/2000 21:55	335,872	74%	86,081	wawa\present\old\
~WRL0357.tmp*	TMP File	25/09/2000 06:08	294,752	90%	28,433	wawa\docs\
nda-schnarch.rtf*	Rich-Text Fo...	26/09/2000 03:43	18,550	80%	3,652	wawa\
netxx695.doc*	Microsoft W...	28/09/2000 10:32	286,909	90%	29,304	wawa\ip\
AtlantaArtcl2.ht...	HTML Docum...	29/09/2000 01:26	34,870	74%	9,083	wawa\

Selected: 1 file, 94KB

Total 1968

start

EXHIBIT G

WinZip - wawa.zip

File Actions View Jobs Options Help

New Open Favorites Add Extract Encrypt View CheckOut Wizard View Style

Name	Type	Modified	Size	Ratio	Packed	Path
ArtdAtInt10925...	HTML Docum...	29/09/2000 01:28	34,399	72%	9,547	wawa\
RtrConfigDDoS.t...	Text Document	29/09/2000 02:12	6,657	54%	3,083	wawa\
~WRL1983.tmp*	TMP File	29/09/2000 06:00	720,384	79%	153,783	wawa\docs\
Copy of netxx.d...	Microsoft W...	29/09/2000 10:07	1,389,056	78%	304,336	wawa\ip\
surgeprotection...	Adobe Acro...	30/09/2000 16:30	125,158	5%	118,692	wawa\
Attack Identifica...	Microsoft W...	02/10/2000 01:13	86,016	81%	16,603	wawa\ip\
TomerQ.doc*	Microsoft W...	02/10/2000 06:41	33,280	87%	4,211	wawa\docs\
TomerQ.zip*	WinZip File	02/10/2000 06:48	4,395	0%	4,407	wawa\docs\
cswsc_wi.htm*	HTML Docum...	05/10/2000 02:05	17,814	65%	6,322	wawa\
Statistical-paten...	Microsoft W...	09/10/2000 05:43	75,264	78%	16,757	wawa\ip\
alpha.doc*	Microsoft W...	09/10/2000 06:04	16,384	81%	3,064	wawa\docs\
netxx.zip*	WinZip File	09/10/2000 07:46	162,601	0%	162,638	wawa\ip\
Stat-par1.doc*	Microsoft W...	09/10/2000 07:54	21,504	88%	2,551	wawa\ip\
netxx3.zip*	WinZip File	09/10/2000 14:47	41,215	0%	41,237	wawa\docs\
1009rev.html*	HTML Docum...	09/10/2000 16:34	44,980	71%	12,889	wawa\docs\
1009cisco.html*	HTML Docum...	09/10/2000 16:36	38,105	71%	10,871	wawa\docs\
report1010.txt*	Text Document	10/10/2000 06:39	1,075	43%	614	wawa\docs\
smallBiz.shtml*	SHTML File	11/10/2000 23:29	83,933	85%	12,472	wawa\
NederlandRTRbi...	HTML Docum...	12/10/2000 14:06	26,789	76%	6,309	wawa\
guide.shtml*	SHTML File	13/10/2000 00:44	39,856	72%	11,245	wawa\
rate.txt*	Text Document	13/10/2000 02:56	5,657	63%	2,075	wawa\docs\
mordi.ppt*	Microsoft Poi...	13/10/2000 04:57	154,624	73%	41,421	wawa\present\old\
cisco_security.Z...	WinZip File	13/10/2000 08:10	1,354,435	0%	1,352,...	wawa\
IOSEssentialsPD...	WinZip File	15/10/2000 08:57	918,492	0%	918,460	wawa\
21.html*	HTML Docum...	15/10/2000 10:27	66,035	67%	21,550	wawa\
ciscoSecurity.html*	HTML Docum...	15/10/2000 10:28	66,035	67%	21,550	wawa\
car.htm*	HTML Docum...	15/10/2000 10:36	47,351	80%	9,566	wawa\
car.txt*	Text Document	15/10/2000 10:44	8,730	60%	3,522	wawa\docs\
ciscoTCPInterse...	HTML Docum...	15/10/2000 13:13	8,735	61%	3,423	wawa\
network.html*	HTML Docum...	16/10/2000 13:04	15,299	57%	6,513	wawa\
TechNews_400...	HTML Docum...	16/10/2000 13:12	61,010	74%	15,677	wawa\
Shernomas.doc*	Microsoft W...	16/10/2000 13:24	5,721	72%	1,627	wawa\ip\
yahooAttack.txt*	Text Document	16/10/2000 23:00	5,766	54%	2,655	wawa\docs\
DammyTarget.p...	PHP3 File	16/10/2000 23:52	10,571	60%	4,226	wawa\docs\
DefenseProposa...	HTML Docum...	17/10/2000 00:48	11,785	62%	4,488	wawa\docs\
newPP.doc*	Microsoft W...	17/10/2000 01:43	20,480	89%	2,314	wawa\docs\

Selected: 1 file, 94KB

Total: 1968

start